arXiv:1704.04118v1 [math.OC] 13 Apr 2017

From Data to Decisions: Distributionally Robust Optimization is Optimal

Bart P.G. Van Parys
Operations Research Center, Massachusetts Institute of Technology, vanparys@mit.edu

Peyman Mohajerin Esfahani
Delft Center for Systems and Control, Technische Universiteit Delft, p.mohajerinesfahani@tudelft.nl

Daniel Kuhn
Risk Analytics and Optimization Chair, Ecole Polytechnique Fédérale de Lausanne, daniel.kuhn@epfl.ch

We study stochastic programs where the decision-maker cannot observe the distribution of the exogenous uncertainties but has access to a finite set of independent samples from this distribution. In this setting, the goal is to find a procedure that transforms the data to an estimate of the expected cost function under the unknown data-generating distribution, *i.e.*, a *predictor*, and an optimizer of the estimated cost function that serves as a near-optimal candidate decision, *i.e.*, a *prescriptor*. As functions of the data, predictors and prescriptors constitute statistical estimators. We propose a meta-optimization problem to find the least conservative predictors and prescriptors subject to constraints on their *out-of-sample disappointment*. The out-of-sample disappointment quantifies the probability that the actual expected cost of the candidate decision under the unknown true distribution exceeds its predicted cost. Leveraging tools from large deviations theory, we prove that this meta-optimization problem admits a unique solution: The best predictor-prescriptor pair is obtained by solving a distributionally robust optimization problem over all distributions within a given relative entropy distance from the empirical distribution of the data.

Key words: Data-Driven Optimization, Distributionally Robust Optimization, Large Deviations Theory, Relative Entropy, Convex Optimization, Observed Fisher Information

1. Introduction

We study static decision problems under uncertainty, where the decision maker cannot observe the probability distribution of the uncertain problem parameters but has access to a finite number of independent samples from this distribution (*e.g.*, a time series). Classical stochastic programming uses this data only indirectly. The data serves as the input for a statistical estimation problem that aims to infer the distribution of the uncertain problem parameters. The estimated distribution then serves as an input for an optimization problem that outputs a near-optimal decision as well as an estimate of the expected cost incurred by this decision. Thus, classical stochastic programming separates the decision-making process into an estimation phase and a subsequent optimization

phase. The estimation method is typically selected with the goal to achieve maximum prediction accuracy but without tailoring it to the optimization problem at hand.

In this paper we develop a method of data-driven stochastic programming that avoids the artificial decoupling of estimation and optimization and that chooses an estimator that adapts to the underlying optimization problem. Specifically, we model data-driven solutions to a stochastic program through a *predictor* and its corresponding *prescriptor*. For any fixed feasible decision, the predictor maps the observable data to an estimate of the decision's expected cost. The predictor, on the other hand, computes a decision that minimizes the cost estimated by the predictor.

The set of all possible predictors and their induced prescriptors is vast. Indeed, there are countless possibilities to estimate the expected costs of a fixed decision from data, *e.g.*, via the popular sample average approximation (Shapiro et al. 2014, Chapter 5), by postulating a parametric model for the exogenous uncertainties and estimating its parameters via maximum likelihood estimation (Dupačová and Wets 1988), or through kernel density estimation (Parpas et al. 2015). Recently, it has become fashionable to construct conservative (pessimistic) estimates of the expected costs via methods of distributionally robust optimization. In this setting, the available data is used to generate an *ambiguity set* that represents a confidence region in the space of probability distributions and contains the unknown data-generating distribution with high probability. The expected cost of a fixed decision under the unknown true distribution is then estimated by the worst-case expectation over all distributions in the ambiguity set. Since the ambiguity set constitutes a confidence region for the unknown true distribution, the worst-case expectation represents an upper confidence bound on the true expected cost. The ambiguity set can be defined, for example, through confidence intervals for the distribution's moments (Delage and Ye 2010). Alternatively, the ambiguity set may contain all distributions that achieve a prescribed level of likelihood (Wang et al. 2016), that pass a statistical hypothesis test (Bertsimas et al. 2014) or that are sufficiently close to a reference distribution with respect to a probability metric such as the Prokhorov metric (Erdoğan and Iyengar 2006), the Wasserstein distance (Pflug and Wozabal 2007, Mohajerin Esfahani and Kuhn 2015, Zhao and Guan 2015), the total variation distance (Sun and Xu 2016) or the L^1-norm (Jiang and Guan 2016a). Ben-Tal et al. (2013) have shown that confidence sets for distributions can also be constructed using ϕ-divergences such as the Pearson divergence, the Burg entropy or the Kullback-Leibler divergence. More recently, Bayraksan and Love (2015) provide a systematic classification of ϕ-divergences and investigate the richness of the corresponding ambiguity sets.

Given the numerous possibilities for constructing predictors from a given dataset, it is easy to loose oversight. In practice, predictors are often selected manually from within a small menu with the goal to meet certain statistical and/or computational requirements. However, there are typically many different predictors that exhibit the desired properties, and there always remains some doubt

as to whether the chosen predictor is best suited for the particular decision problem at hand. In this paper we propose a principled approach to data-driven stochastic programming by solving a meta-optimization problem over a rich class of predictor-prescriptor-pairs including, among others, all examples reviewed above. This meta-optimization problem aims to find the least conservative (*i.e.*, pointwise smallest) prescriptor whose *out-of-sample disappointment* decays at a prescribed exponential rate r as the sample size tends to infinity—irrespective of the true data-generating distribution. The out-of-sample disappointment quantifies the probability that the *actual* expected cost of the prescriptor exceeds its *predicted* cost. Put differently, it represents the probability that the predicted cost of a candidate decision is over-optimistic and leads to disappointment in out-of-sample tests. Thus, the proposed meta-optimization problem tries to identify the predictor-prescriptor pairs that overestimate the expected out-of-sample costs by the least amount possible without risking disappointment under *any* thinkable data-generating distribution.

Our main results can be summarized as follows.

• By leveraging Sanov's theorem from large deviations theory, we prove that the meta-optimization problem admits a unique optimal solution for any given stochastic program.

• We show that the optimal data-driven predictor estimates the expected costs under the unknown true distribution by a worst-case expectation over all distributions within a given relative entropy distance from the empirical distribution of the data. This suggests that, among all possible data-driven solutions, a distributionally robust approach based on a relative entropy ambiguity set is optimal. This is perhaps surprising because the meta-optimization problem does not impose any structure on the predictors, which are generic functions of the data. In particular, there is no requirement forcing predictors to admit a distributionally robust interpretation.

• In marked contrast to the existing work on data-driven distributionally robust optimization, our relative entropy ambiguity set does *not* play the role of a confidence region that contains the unknown data-generating distribution with a prescribed level of probability. Instead, the radius of the relative entropy ambiguity set coincides with the desired exponential decay rate r of the out-of-sample disappointment imposed by the meta-optimization problem.

• The conservatism of the optimal predictor, defined as its deviation from the expected cost under the empirical distribution, can be viewed as the price to combat out-of-sample disappointment. We show that, to leading order in r, the conservatism is proportional to $\sqrt{r}$ multiplied by the empirical standard deviation of the costs. Thus, the price to combat out-of-sample disappointment depends on the observed dataset, indicating that some observations are more valuable than others.

To our best knowledge, we are the first to recognize the optimality of distributionally robust optimization in its ability to transform data to predictors and prescriptors. The optimal distributionally robust predictor identified in this paper can be evaluated by solving a tractable convex

optimization problem. Under standard convexity assumptions about the feasible set and the cost function of the stochastic program, the corresponding optimal prescriptor can also be evaluated in polynomial time. Although perhaps desirable, the tractability and distributionally robust nature of the optimal predictor-prescriptor pair are not dictated *ex ante* but emerge naturally.

Relative entropy ambiguity sets have already attracted considerable interest in distributionally robust optimization (Ben-Tal et al. 2013, Calafiore 2007, Hu and Hong 2013, Lam 2016b, Wang et al. 2016). Note, however, that the relative entropy constitutes an *asymmetric* distance measure between two distributions. The asymmetry implies, among others, that the first distribution must be absolutely continuous to the second one but not *vice versa*. Thus, ambiguity sets can be constructed in two different ways by designating the reference distribution either as the first or as the *second* argument of the relative entropy. All papers listed above favor the second option, and thus the emerging ambiguity sets contain only distributions that are absolutely continuous to the reference distribution. Maybe surprisingly, the optimal predictor resulting from our meta-optimization problem uses the reference distribution as the *first* argument of the relative entropy instead. Thus, the reference distribution is absolutely continuous to every distribution in the emerging ambiguity set. Relative entropy balls of this kind have previously been studied in (Gupta 2015, Lam 2016a).

Adopting a Bayesian perspective, Gupta (2015) determines the smallest ambiguity sets that contain the unknown data-generating distribution with a prescribed level of confidence as the sample size tends to infinity. Both Pearson divergence and relative entropy ambiguity sets with properly scaled radii are optimal in this setting. In the terminology of the present paper, Gupta (2015) thus restricts attention to the subclass of distributionally robust predictors and operates with an asymptotic notion of optimality. The meta-optimization problem proposed here entails a stronger notion of optimality, under which the distributionally robust predictor with relative entropy ambiguity set emerges as the unique optimizer. Lam (2016a) also seeks distributionally robust predictors that trade conservatism for out-of-sample performance. He studies the probability that the estimated expected cost function dominates the actual expected cost function uniformly across all decisions, and he calls a predictor optimal if this probability is asymptotically equal to a prescribed confidence level. Using the empirical likelihood theorem of Owen (1988), he shows that Pearson divergence and relative entropy ambiguity sets with properly scaled radii are optimal in this sense. This notion of optimality has again an asymptotic flavor because it refers to decreasing sequences of ambiguity sets that converge to a singleton, and it admits multiple optimizers.

The rest of the paper unfolds as follows. Section 2 provides a formal introduction to data-driven stochastic programming and develops the meta-optimization problem for identifying the best predictor-prescriptor-pair. Section 3 reviews weak and strong large deviation principles, which are then used in Section 4 to determine the unique optimal solution of the meta-optimization problem. The price to combat out-of-sample disappointment is investigated in Section 5.

Notation: The natural logarithm of $p \in \Re_+$ is denoted by $\log(p)$, where we use the conventions $0\log(0/p) = 0$ for any $p \geq 0$ and $p'\log(p'/0) = \infty$ for any $p' > 0$. The set of symmetric matrices in $\Re^{d\times d}$ is denoted by $\mathbb{S}^d$. For $F, H \in \mathbb{S}^d$ the relation $F \succeq H$ ($F \succ H$) indicates that $F - H$ is positive semidefinite (positive definite). The symbol $\mathbb{I}$ stands for the identity matrix in $\Re^{d\times d}$. For $z \in \Re^d$ we denote by $\mathrm{diag}(z) \in \Re^{d\times d}$ the diagonal matrix with z on the main diagonal. A function $f : \mathcal{P} \to X$ from $\mathcal{P} \subseteq \Re^d$ to $X \subseteq \Re^n$ is called quasi-continuous at $\mathbb{P} \in \mathcal{P}$ if for every $\epsilon > 0$ and neighborhood $U \subseteq \mathcal{P}$ of $\mathbb{P}$ there is a non-empty open set $V \subseteq U$ with $|f(\mathbb{P}) - f(\mathbb{Q})| \leq \epsilon$ for all $\mathbb{Q} \in V$. Note that V does not necessarily contain $\mathbb{P}$. For any logical statement $\mathcal{E}$, the indicator function $\mathbb{1}_{\mathcal{E}}$ evaluates to 1 if $\mathcal{E}$ is true and to 0 otherwise.

2. Data-driven stochastic programming

Stochastic programming is a powerful modeling paradigm for taking informed decisions in an uncertain environment. A generic single-stage stochastic program can be represented as

$$\underset{x\in X}{\text{minimize}}\ \mathbb{E}_{\mathbb{P}^\star}\left[\gamma(x,\xi)\right]. \tag{1}$$

Here, the goal is to minimize the expected value of a cost function $\gamma(x,\xi) \in \Re$, which depends both on a decision variable $x \in X$ and a random parameter $\xi \in \Xi$ governed by a probability distribution $\mathbb{P}^\star$. Below we will assume that the cost $\gamma(x,\xi)$ is continuous in x for every fixed $\xi \in \Xi$, the feasible set $X \subseteq \Re^n$ is compact, and $\Xi = \{1,\ldots,d\}$ is finite. Thus, ξ has d distinct scenarios that are represented—without loss of generality—by the integers $1,\ldots,d$. A wide spectrum of decision problems can be cast as instances of (1). Shapiro et al. (2014) point out, for example, that (1) can be viewed as the first stage of a two-stage stochastic program, where the cost function $\gamma(x,\xi)$ embodies the optimal value of a subordinate second-stage problem. Alternatively, problem (1) may also be interpreted as a generic learning problem in the spirit of statistical learning theory.

In the following, we distinguish the *prediction problem*, which merely aims to predict the expected cost associated with a fixed decision x, and the *prescription problem*, which seeks to identify a decision $x^\star$ that minimizes the expected cost across all $x \in X$.

Any attempt to solve the prescription problem seems futile unless there is a procedure for solving the corresponding prediction problem. The generic prediction problem is closely related to what Le Maître and Knio (2010) call an uncertainty quantification problem and is therefore of prime interest in its own right. Throughout the rest of the paper, we thus analyze prediction and prescription problems on equal footing.

In the what follows we formalize the notion of a data-driven solution to the prescription and prediction problems, respectively. Furthermore, we introduce the basic assumptions as well as the notation used throughout the remainder of the paper.

2.1. Data-driven predictors and prescriptors

Unfortunately, the distribution $\mathbb{P}^\star$ of ξ is hardly ever observable but must be estimated from a training dataset, that is, a finite (possibly small) number of independent samples from $\mathbb{P}^\star$. Thus, we lack essential information to evaluate the expected cost of any fixed decision and—*a fortiori*—to solve the stochastic program (1). The standard approach to overcome this deficiency is to approximate $\mathbb{P}^\star$ with a parametric or non-parametric estimate $\hat{\mathbb{P}}$ inferred from the samples and to minimize the expected cost under $\hat{\mathbb{P}}$ instead of the true expected cost under $\mathbb{P}^\star$. However, if we calibrate a stochastic program to a training data set and evaluate its optimal decision on a test data set, then the resulting test performance is often disappointing—even if the two datasets are sampled independently from $\mathbb{P}^\star$. This phenomenon has been observed in many different contexts. It is particularly pronounced in finance, where Michaud (1989) refers to it as the 'error maximization effect' of portfolio optimization, and in statistics or machine learning, where it is known as 'overfitting'. In decision analysis, Smith and Winkler (2006) refer to it as the 'optimizer's curse'. Thus, when working with data instead of exact probability distributions, one should safeguard against solutions that display promising in-sample performance but lead to out-of-sample disappointment.

Initially the distribution $\mathbb{P}^\star$ is only known to belong to the probability simplex $\mathcal{P} = \{\mathbb{P} \in \Re^d_+ : \sum_{i\in\Xi} \mathbb{P}(i) = 1\}$. Over time, however, independent samples ξ_t, $t \in \mathbb{N}$, from $\mathbb{P}^\star$ are revealed to the decision maker that provide increasingly reliable statistical information about $\mathbb{P}^\star$.

Any $\mathbb{P} \in \mathcal{P}$ encodes a possible probabilistic model for the data process. Thus, by slight abuse of terminology, we will henceforth refer to the distributions $\mathbb{P} \in \mathcal{P}$ as models and to $\mathcal{P}$ as the model class. Evidently, the true model $\mathbb{P}^\star$ is an (albeit unknown) element of $\mathcal{P}$. Next, we introduce parametric predictors and prescriptors corresponding to the stochastic program (1), where the true unknown distribution $\mathbb{P}^\star$ is replaced with a hypothetical model $\mathbb{P} \in \mathcal{P}$.

Definition 1 (Parametric predictors and prescriptors). For any fixed model $\mathbb{P} \in \mathcal{P}$, we define the *parametric predictor* $c(x, \mathbb{P}) = \mathbb{E}_{\mathbb{P}}[\gamma(x, \xi)] = \sum_{i\in\Xi} \mathbb{P}(i)\, \gamma(x, i)$ as the expected cost of a given decision $x \in X$ and the *parametric prescriptor* $x^\star(\mathbb{P}) \in \arg\min_{x\in X} c(x, \mathbb{P})$ as a decision that minimizes $c(x, \mathbb{P})$ over $x \in X$.

Note that parametric predictor $c(x, \mathbb{P})$ is jointly continuous in x and $\mathbb{P}$ because Ξ is finite and $\gamma(x, \xi)$ is continuous in x for every fixed $\xi \in \Xi$. The continuity of $c(x, \mathbb{P})$ then guarantees via the compactness of X that the parametric prescriptor $x^\star(\mathbb{P})$ exists for every model $\mathbb{P} \in \mathcal{P}$. In view of Definition 1, the stochastic program (1) can be identified with the *prescription problem* of computing $x^\star(\mathbb{P}^\star)$. Similarly, the evaluation of the expected cost of a given decision $x \in X$ in (1) can be identified with the *prediction problem* of computing $c(x, \mathbb{P}^\star)$. These prediction and prescription problems cannot be solved, however, as they depend on the unknown true model $\mathbb{P}^\star$.

If one has only access to a finite set $\{\xi_t\}_{t=1}^T$ of independent samples from $\mathbb{P}^\star$ instead of $\mathbb{P}^\star$ itself, then it may be useful to construct an empirical estimator for $\mathbb{P}^\star$.

Definition 2 (Empirical Distribution). The empirical distribution $\hat{\mathbb{P}}_T$ corresponding to the sample path $\{\xi_t\}_{t=1}^T$ of length T is defined through

$$\hat{\mathbb{P}}_T(i) = \frac{1}{T}\sum_{t=1}^{T} \mathbb{1}_{\xi_t=i} \quad \forall i \in \Xi.$$

Note that $\hat{\mathbb{P}}_T$ can be viewed as the vector of empirical state frequencies. Indeed, its i^{th} entry records the proportion of time that the sample path spends in state i. As the samples are drawn independently, the state frequencies capture all useful statistical information about $\mathbb{P}^\star$ that can possibly be extracted from a given sample path. Note also that $\hat{\mathbb{P}}$ is in fact the maximum likelihood estimator of $\mathbb{P}^\star$. In the following, we will therefore approximate the unknown predictor $c(x, \mathbb{P}^\star)$ as well as the unknown prescriptor $x^\star(\mathbb{P}^\star)$ by suitable functions of the empirical distribution $\hat{\mathbb{P}}_T$.

Definition 3 (Data-driven predictors and prescriptors). A continuous function $\hat{c} : X \times \mathcal{P} \to \Re$ is called a *data-driven predictor* if $\hat{c}(x, \hat{\mathbb{P}}_T)$ is used as an approximation for $c(x, \mathbb{P}^\star)$. A quasi-continuous function $\hat{x} : \mathcal{P} \to X$ is called a *data-driven prescriptor* if there exists a data-driven predictor $\hat{c}$ with

$$\hat{x}(\mathbb{P}') \in \arg\min_{x \in X} \hat{c}(x, \mathbb{P}')$$

for all possible estimator realizations $\mathbb{P}' \in \mathcal{P}$, and $\hat{x}(\hat{\mathbb{P}}_T)$ is used as an approximation for $x^\star(\mathbb{P}^\star)$.

Theorem 5.3 in (Matejdes 1987) guarantees that every data-driven predictor $\hat{c}$ induces a data-driven prescriptor because X is compact, while the 'arg min' mapping is non-empty-valued and upper semicontinuous. Note that upper semicontinuity follows from Berge's maximum theorem (Berge 1963, pp. 115–116), which applies because $\hat{c}$ is continuous and X is independent of $\mathbb{P}'$. We remark that the set of points where a quasi-continuous prescriptor is continuous is dense in $\mathcal{P}$ (Bledsoe 1952). Thus, data-driven prescriptors in the sense of Definition 3 are 'mostly' continuous.

Example 1 (Sample average predictor). The parametric predictor c introduced in Definition 1 constitutes a simple data-driven predictor $\hat{c} = c$, that is, $c(x, \hat{\mathbb{P}}_T)$ can readily be used as a naïve approximation for $c(x, \mathbb{P}^\star)$. Note that the parametric predictor c is indeed continuous as desired. By the definition of the empirical estimator, this naïve predictor approximates $c(x, \mathbb{P}^\star)$ with

$$c(x, \hat{\mathbb{P}}_T) = \frac{1}{T}\sum_{t=1}^{T} \gamma(x, \xi_t),$$

which is readily recognized as the popular sample average approximation.

2.2. Optimizing over all data-driven predictors and prescriptors

The estimates $\hat{c}(x, \hat{\mathbb{P}}_T)$ and $\hat{x}(\hat{\mathbb{P}}_T)$ inherit the randomness from the empirical estimator $\hat{\mathbb{P}}_T$, which is constructed from the (random) samples $\{\xi_t\}_{t=1}^T$. Note that the prediction and prescription problems are naturally interpreted as instances of statistical estimation problems. Indeed, data-driven

prediction aims to estimate the expected cost $c(x, \mathbb{P}^\star)$ from data. Standard statistical estimation theory would typically endeavor to find a data-driven predictor $\hat{c}$ that (approximately) minimizes the mean squared error

$$\mathbb{E}\left[|c(x, \mathbb{P}^\star) - \hat{c}(x, \hat{\mathbb{P}}_T)|^2\right]$$

over some appropriately chosen class of predictors $\hat{c}$, where the expectation is taken with respect to the distribution $(\mathbb{P}^\star)^\infty$ governing the sample path and the empirical estimator. The mean squared error penalizes the mismatch between the actual cost $c(x, \mathbb{P}^\star)$ and its estimator $\hat{c}(x, \hat{\mathbb{P}}_T)$. Events in which we are left disappointed ($c(x, \mathbb{P}^\star) > \hat{c}(x, \hat{\mathbb{P}}_T)$) are not treated differently from positive surprises ($c(x, \mathbb{P}^\star) < \hat{c}(x, \hat{\mathbb{P}}_T)$). In a decision-making context where the goal is to minimize costs, however, disappointments (underestimated costs) are more harmful than positive surprises (overestimated costs). While statisticians strive for accuracy by minimizing a symmetric estimation error, decision makers endeavor to limit the one-sided prediction disappointment.

Definition 4 (Out-of-sample disappointment). For any data-driven predictor $\hat{c}$ the probability

$$\mathbb{P}^\infty\left(c(x, \mathbb{P}) > \hat{c}(x, \hat{\mathbb{P}}_T)\right) \tag{2a}$$

is referred to as the out-of-sample prediction disappointment of $x \in X$ under model $\mathbb{P} \in \mathcal{P}$. Similarly, for any data-driven prescriptor $\hat{x}$ induced by a data-driven predictor $\hat{c}$ the probability

$$\mathbb{P}^\infty\left(c(\hat{x}(\hat{\mathbb{P}}_T), \mathbb{P}) > \hat{c}(\hat{x}(\hat{\mathbb{P}}_T), \hat{\mathbb{P}}_T)\right) \tag{2b}$$

is termed the out-of-sample prescription disappointment under model $\mathbb{P} \in \mathcal{P}$.

The out-of-sample prediction disappointment quantifies the probability (with respect to the sample path distribution $\mathbb{P}^\infty$ under some model $\mathbb{P} \in \mathcal{P}$) that the expected cost $c(x, \mathbb{P})$ of a fixed decision x exceeds the predicted cost $\hat{c}(x, \hat{\mathbb{P}}_T)$. Thus, the out-of-sample prediction disappointment is independent of the actual realization of the empirical estimator $\hat{\mathbb{P}}_T$ but depends on the hypothesized model $\mathbb{P}$. A similar statement holds for the out-of-sample prescription disappointment.

The main objective of this paper is to construct attractive data-driven predictors and prescriptors, which are optimal in a sense to be made precise below. We first develop a notion of optimality for data-driven predictors and extended it later to data-driven prescriptors. As indicated above, a crucial requirement for any data-driven predictor is that it must limit the out-of-sample disappointment. This informal requirement can be operationalized either in an asymptotic sense or in a finite sample sense.

(i) **Asymptotic guarantee:** As T grows, the out-of-sample prediction disappointment (2a) decays exponentially at a rate at least equal to $r > 0$ up to first order in the exponent, that is,

$$\limsup_{T\to\infty} \frac{1}{T} \log \mathbb{P}^\infty\left(c(x, \mathbb{P}) > \hat{c}(x, \hat{\mathbb{P}}_T)\right) \leq -r \quad \forall x \in X,\ \mathbb{P} \in \mathcal{P}. \tag{3}$$

(ii) **Finite sample guarantee:** For every fixed T, the out-of-sample prediction disappointment (2a) is bounded above by a *known* function $g(T)$ that decays exponentially at rate at least equal to $r > 0$ to first order in the exponent, that is,

$$\mathbb{P}^\infty \left(c(x, \mathbb{P}) > \hat{c}(x, \hat{\mathbb{P}}_T) \right) \leq g(T) \quad \forall x \in X,\ \mathbb{P} \in \mathcal{P},\ T \in \mathbb{N}, \tag{4}$$

where $\limsup_{T\to\infty} \frac{1}{T} \log g(T) \leq -r$.

The inequalities (3) and (4) are imposed across all models $\mathbb{P} \in \mathcal{P}$. This ensures that they are satisfied under the true model $\mathbb{P}^\star$, which is only known to reside within $\mathcal{P}$. By requiring the inequalities to hold for all $x \in X$, we further ensure that the out-of-sample prediction disappointment is eventually small irrespective of the chosen decision. Note that the finite sample guarantee (4) is sufficient but not necessary for the asymptotic guarantee (3). Knowing the finite sample bounds $g(T)$ has the advantage, amongst others, that one can determine the sample complexity

$$\min \left\{ T_0 \in \mathbb{N} : g(T) \leq \beta,\ \forall T \geq T_0 \right\},$$

that is, the minimum number of samples needed to certify that the out-of-sample prediction disappointment does not exceed a prescribed significance level $\beta \in [0, 1]$.

At first sight the requirements (3) and (4) may seem unduly restrictive, and the existence of data-driven predictors with exponentially decaying out-of-sample disappointment may be questioned. Indeed, the popular sample average predictor from Example 1 violates these requirements.

Example 2 (Large out-of-sample disappointment). Set the cost function to $\gamma(x, \xi) = \xi$. In this case, the sample average predictor approximates the expected cost $c(x, \mathbb{P}) = \sum_{i\in\Xi} i\mathbb{P}(i)$ by its sample mean $c(x, \hat{\mathbb{P}}_T) = \frac{1}{T} \sum_{t=1}^T \xi_t$. As the sample size T tends to infinity, the central limit theorem implies that

$$\sqrt{T}[c(x, \hat{\mathbb{P}}_T) - c(x, \mathbb{P})]$$

converges in law to a normal distribution with mean 0 and variance $\sum_{i\in\Xi} i^2\mathbb{P}(i) - (\sum_{i\in\Xi} i\mathbb{P}(i))^2$. Thus,

$$\lim_{T\to\infty} \mathbb{P}^\infty \left(c(x, \mathbb{P}) > \hat{c}(x, \hat{\mathbb{P}}_T) \right) = \lim_{T\to\infty} \mathbb{P}^\infty \left(\sqrt{T} \left(\hat{c}(x, \hat{\mathbb{P}}_T) - c(x, \mathbb{P}) \right) < 0 \right) = \frac{1}{2},$$

which means that the out-of-sample prediction disappointment remains large for all sample sizes. The sample average predictor hence violates the asymptotic guarantee (3) and—*a fortriori*—also the stronger finite sample guarantee (4).

Example 2 suggests that the out-of-sample disappointment of a predictor $\hat{c}$ cannot be expected to decay at an exponential rate unless $\hat{c}$ is *conservative*, that is, unless $\hat{c}(x, \mathbb{P}') > c(x, \mathbb{P}')$ for all decisions $x \in X$ and estimator realizations $\mathbb{P}' \in \mathcal{P}$. If the predictor is conservative, then—maybe surprisingly—an exponential decay of the prediction disappointment is to be expected under rather

generic conditions because the empirical estimator $\hat{\mathbb{P}}_T$ obeys a strong law of large numbers under $\mathbb{P}^\infty$ (meaning that $\hat{\mathbb{P}}_T$ converges $\mathbb{P}^\infty$-almost surely to $\mathbb{P}$). In fact, asymptotic guarantees of the type (3) hold whenever the empirical distribution $\hat{\mathbb{P}}_T$ satisfies a *weak large deviation principle*, while finite sample guarantees of the type (4) hold when $\hat{\mathbb{P}}_T$ satisfies a *strong large deviation principle*. As will be shown in Section 3, the empirical distribution satisfies a strong large deviation principle.

For ease of exposition, we henceforth denote by $\mathcal{C}$ the set of all data-driven predictors, that is, all continuous functions that map $X \times \mathcal{P}$ to the reals. Moreover, we introduce a partial order $\preceq_{\mathcal{C}}$ on $\mathcal{C}$ defined through

$$\hat{c}_1 \preceq_{\mathcal{C}} \hat{c}_2 \quad \Longleftrightarrow \quad \hat{c}_1(x, \mathbb{P}') \leq \hat{c}_2(x, \mathbb{P}') \quad \forall x \in X,\, \mathbb{P}' \in \mathcal{P}$$

for any $\hat{c}_1, \hat{c}_2 \in \mathcal{C}$. Thus, $\hat{c}_1 \preceq_{\mathcal{C}} \hat{c}_2$ means that $\hat{c}_1$ is (weakly) less conservative than $\hat{c}_2$. The problem of finding the least conservative predictor among all data-driven predictors whose out-of-sample disappointment decays at rate at least $r > 0$ can thus be formalized as the following *vector optimization problem*.

$$\begin{array}{ll} \underset{\hat{c} \in \mathcal{C}}{\text{minimize}}_{\preceq_{\mathcal{C}}} & \hat{c} \\ \text{subject to} & \limsup\limits_{T \to \infty} \dfrac{1}{T} \log \mathbb{P}^\infty \left(c(x, \mathbb{P}) > \hat{c}(x, \hat{\mathbb{P}}_T) \right) \leq -r \quad \forall x \in X,\, \mathbb{P} \in \mathcal{P} \end{array} \tag{5}$$

We highlight that the minimization in (5) is understood with respect to the partial order $\preceq_{\mathcal{C}}$. Thus, the relation $\hat{c}_1 \preceq_{\mathcal{C}} \hat{c}_2$ between two feasible decision means that $\hat{c}_1$ is weakly preferred to $\hat{c}_2$. However, not all pairs of feasible decisions are comparable, that is, it is possible that both $\hat{c}_1 \not\preceq_{\mathcal{C}} \hat{c}_2$ and $\hat{c}_2 \not\preceq_{\mathcal{C}} \hat{c}_1$. A predictor $\hat{c}^\star$ is a *strongly* optimal solution for (5) if it is feasible and weakly preferred to every other feasible solution (*i.e.*, every $\hat{c} \neq \hat{c}^\star$ feasible in (5) satisfies $\hat{c}^\star \preceq_{\mathcal{C}} \hat{c}$). Similarly, $\hat{c}^\star$ is a *weakly* optimal solution for (5) if it is feasible and if every other solution preferred to $\hat{c}^\star$ is infeasible (*i.e.*, every $\hat{c} \neq \hat{c}^\star$ with $\hat{c} \preceq_{\mathcal{C}} \hat{c}^\star$ is infeasible in (5)).

We are now ready to construct a meta-optimization problem akin to (5), which enables us to identify the best prescriptor. To this end, we henceforth denote by $\mathcal{X}$ the set of all data-driven predictor-prescriptor-pairs $(\hat{c}, \hat{x})$, where $\hat{c} \in \mathcal{C}$, and $\hat{x}$ is a prescriptor induced by $\hat{c}$ as per Definition 3. Moreover, we equip $\mathcal{X}$ with a partial order $\preceq_{\mathcal{X}}$, which is defined through

$$(\hat{c}_1, \hat{x}_1) \preceq_{\mathcal{X}} (\hat{c}_2, \hat{x}_2) \quad \Longleftrightarrow \quad \hat{c}_1(\hat{x}_1(\mathbb{P}'), \mathbb{P}') \leq \hat{c}_2(\hat{x}_2(\mathbb{P}'), \mathbb{P}') \quad \forall \mathbb{P}' \in \mathcal{P}.$$

Note that $\hat{c}_1 \preceq_{\mathcal{C}} \hat{c}_2$ actually implies $(\hat{c}_1, \hat{x}_1) \preceq_{\mathcal{X}} (\hat{c}_2, \hat{x}_2)$ but not vice versa. The problem of finding the least conservative predictor-prescriptor-pair whose out-of-sample prescription disappointment decays at rate at least $r > 0$ can now be formalized as the following vector optimization problem.

$$\begin{array}{ll} \underset{(\hat{c}, \hat{x}) \in \mathcal{X}}{\text{minimize}}_{\preceq_{\mathcal{X}}} & (\hat{c}, \hat{x}) \\ \text{subject to} & \limsup\limits_{T \to \infty} \dfrac{1}{T} \log \mathbb{P}^\infty \left(c(\hat{x}(\hat{\mathbb{P}}_T), \mathbb{P}) > \hat{c}(\hat{x}(\hat{\mathbb{P}}_T), \hat{\mathbb{P}}_T) \right) \leq -r \quad \forall \mathbb{P} \in \mathcal{P} \end{array} \tag{6}$$

Generic vector optimization problems typically only admit weak solutions. In Section 4 we will show, however, that (5) as well as (6) admit strong solutions in closed form. In fact, we will show that these closed-form solutions have a natural interpretation as the solutions of convex distributionally robust optimization problems.

3. Large deviation principles

Large deviations theory provides bounds on the exact exponential rate at which the probabilities of atypical estimator realizations decay under a model $\mathbb{P}$ as the sample size T tends to infinity. These bounds are expressed in terms of the relative entropy of $\hat{\mathbb{P}}_T$ with respect to $\mathbb{P}$.

Definition 5 (Relative entropy). The relative entropy of an estimator realization $\mathbb{P}' \in \mathcal{P}$ with respect to a model $\mathbb{P} \in \mathcal{P}$ is defined as

$$\mathrm{I}(\mathbb{P}', \mathbb{P}) = \sum_{i \in \Xi} \mathbb{P}'(i) \log\left(\frac{\mathbb{P}'(i)}{\mathbb{P}(i)}\right),$$

where we use the conventions $0 \log(0/p) = 0$ for any $p \geq 0$ and $p' \log(p'/0) = \infty$ for any $p' > 0$.

The relative entropy is also known as information for discrimination, cross-entropy, information gain or Kullback-Leibler divergence (Kullback and Leibler 1951). The following proposition summarizes the key properties of the relative entropy relevant for this paper.

Proposition 1 **(Relative entropy)**. *The relative entropy enjoys the following properties:*

(i) **Information inequality:** $\mathrm{I}(\mathbb{P}', \mathbb{P}) \geq 0$ *for all* $\mathbb{P}, \mathbb{P}' \in \mathcal{P}$*, while* $\mathrm{I}(\mathbb{P}', \mathbb{P}) = 0$ *if and only if* $\mathbb{P}' = \mathbb{P}$.

(ii) **Convexity:** *For all pairs* $(\mathbb{P}'_1, \mathbb{P}_1), (\mathbb{P}'_2, \mathbb{P}_2) \in \mathcal{P} \times \mathcal{P}$ *and* $\lambda \in [0, 1]$ *we have*

$$\mathrm{I}((1-\lambda)\mathbb{P}'_1 + \lambda \mathbb{P}'_2, (1-\lambda)\mathbb{P}_1 + \lambda \mathbb{P}_2) \leq (1-\lambda)\mathrm{I}(\mathbb{P}'_1, \mathbb{P}_1) + \lambda \mathrm{I}(\mathbb{P}'_2, \mathbb{P}_2).$$

(iii) **Lower semicontinuity** $\mathrm{I}(\mathbb{P}', \mathbb{P}) \geq 0$ *is lower semicontinuous in* $(\mathbb{P}', \mathbb{P}) \in \mathcal{P} \times \mathcal{P}$.

Proof. Assertions (i) and (ii) follow from Theorems 2.6.3 and 2.7.2 in Cover and Thomas (2006), respectively, while assertion (iii) follows directly from the definition of the relative entropy and our standard conventions regarding the natural logarithm. □

We now show that the empirical estimators satisfy a weak *large deviation principle* (LDP). This result follows immediately from a finite version of Sanov's classical theorem. A textbook proof using the so-called method of types can be found in Cover and Thomas (2006, Theorem 12.4.1). As the proof is illuminating and to keep this paper self-contained, we sketch the proof in Appendix A.

Theorem 1 **(Weak LDP)**. *If the samples* $\{\xi_t\}_{t \in \mathbb{N}}$ *are drawn independently from some* $\mathbb{P} \in \mathcal{P}$*, then for every Borel set* $\mathcal{D} \subseteq \mathcal{P}$ *the sequence of empirical distributions* $\{\hat{\mathbb{P}}_T\}_{T \in \mathcal{P}}$ *satisfies*

$$\limsup_{T \to \infty} \frac{1}{T} \log \mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{D}) \leq - \inf_{\mathbb{P}' \in \mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P}). \tag{7a}$$

If additionally $\mathbb{P} > 0$*, then for every Borel set* $\mathcal{D} \subseteq \mathcal{P}$ *we have*

$$\liminf_{T\to\infty} \ \frac{1}{T} \log \mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{D}) \geq - \inf_{\mathbb{P}' \in \operatorname{int} \mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P}). \tag{7b}$$

Note that the inequality (7a) provides an *upper* LDP bound on the exponential rate at which the probability of the event $\hat{\mathbb{P}}_T \in \mathcal{D}$ decays under model $\mathbb{P}$. This upper bound is expressed in terms of a convex optimization problem that minimizes the relative entropy of $\mathbb{P}'$ with respect to $\mathbb{P}$ across all estimator realizations $\mathbb{P}'$ within $\mathcal{D}$. Similarly, (7b) offers a *lower* LDP bound on the decay rate. Note that in (7b) the relative entropy is minimized over the *interior* of $\mathcal{D}$ instead of $\mathcal{D}$.

If the data-generating model $\mathbb{P}$ itself belongs to $\mathcal{D}$, then $\inf_{\mathbb{P}' \in \mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P}) = \mathrm{I}(\mathbb{P}, \mathbb{P}) = 0$, which leads to the trivial upper bound $\mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{D}) \leq 1$. On the other hand, if $\mathcal{D}$ has empty interior (*e.g.*, if $\mathcal{D} = \{\mathbb{P}\}$ is a singleton containing only the true model), then $\inf_{\mathbb{P}' \in \operatorname{int} \mathcal{D}} I(\mathbb{P}', \mathbb{P}) = \infty$, which leads to the trivial lower bound $\mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{D}) \geq 0$. Non-trivial bounds are obtained if $\mathbb{P} \notin \mathcal{D}$ and $\operatorname{int} \mathcal{D} \neq \emptyset$. In these cases the relative entropy bounds the exponential rate at which the probability of the atypical event $\hat{\mathbb{P}}_T \in \mathcal{D}$ decays with T. For some sets $\mathcal{D}$ this rate of decay is precisely determined by the relative entropy. Specifically, a Borel set $\mathcal{D} \subseteq \mathcal{P}$ is called I-continuous under model $\mathbb{P}$ if

$$\inf_{\mathbb{P}' \in \operatorname{int} \mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P}) = \inf_{\mathbb{P}' \in \mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P}).$$

Clearly, every open set $\mathcal{D} \subseteq \mathcal{P}$ is I-continuous under any model $\mathbb{P}$. As the relative entropy is continuous in $\mathbb{P}'$ whenever $\mathbb{P} > 0$, moreover, any Borel set $\mathcal{D} \subseteq \mathcal{P}$ with $\mathcal{D} \subseteq \operatorname{cl}(\operatorname{int}(\mathcal{D})$ is I-continuous under $\mathbb{P}$ if $\mathbb{P} > 0$. The LDP (7) implies that for large T the probability of an I-continuous set $\mathcal{D}$ decays at rate $\inf_{\mathbb{P}' \in \mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P})$ under model $\mathbb{P}$ to first order in the exponent, that is, we have

$$\mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{D}) = e^{-T \inf_{\mathbb{P}' \in \mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P}) + o(T)}. \tag{8}$$

If we interpret the relative entropy $\mathrm{I}(\mathbb{P}', \mathbb{P})$ as the distance of $\mathbb{P}$ from $\mathbb{P}'$, then the decay rate of $\mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{D})$ coincides with the distance of the model $\mathbb{P}$ from the atypical event set $\mathcal{D}$; see Figure 1. Moreover, if $\mathcal{D}$ is I-continuous under $\mathbb{P}$, then (8) implies that $\mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{D}) \leq \beta$ whenever

$$T \gtrsim \frac{1}{r} \cdot \log\left(\frac{1}{\beta}\right),$$

where $r = \inf_{\mathbb{P}' \in \mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P})$ is the I-distance from $\mathbb{P}$ to the set $\mathcal{D}$, and $\beta \in (0, 1)$ is a prescribed significance level.

The weak LDP of Theorem 1 provides only *asymptotic* bounds on the decay rates of atypical events. However, one can also establish a *strong* LDP, which offers *finite sample guarantees*. Most results of this paper, however, are based on the weak LDP of Theorem 1.

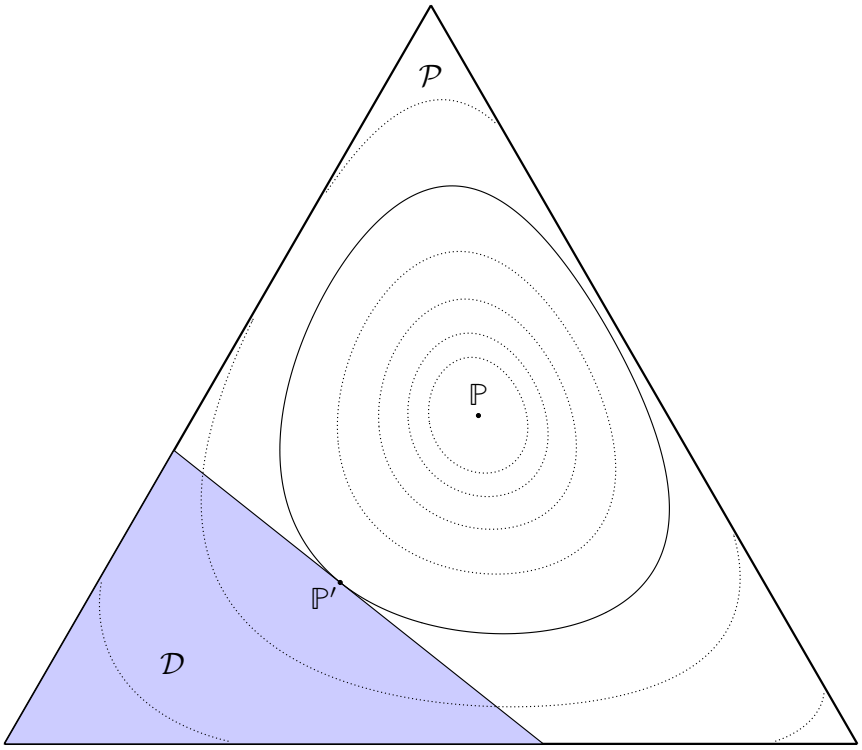

Figure 1 Visualization of the LDP (7). If $\mathcal{D} \subseteq \mathcal{P}$ is I-continuous and $\mathbb{P} \notin \mathcal{D}$, then the probability $\mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{D})$ decays at the exponential rate $\inf_{\mathbb{P}' \in \mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P})$, which can be viewed as the relative entropy distance of $\mathbb{P}$ from $\mathcal{D}$.

THEOREM 2 **(Strong LDP)**. *If the samples $\{\xi_t\}_{t\in\mathbb{N}}$ are drawn independently from some $\mathbb{P} \in \mathcal{P}$, then for every Borel set $\mathcal{D} \subseteq \mathcal{P}$ the sequence of empirical distributions $\{\hat{\mathbb{P}}_T\}_{T\in\mathcal{P}}$ satisfies*

$$\mathbb{P}^\infty \left(\hat{\mathbb{P}}_T \in \mathcal{D} \right) \leq (T+1)^d e^{-T \inf_{\mathbb{P}' \in \mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P})} \quad \forall T \in \mathbb{N}. \tag{9}$$

Proof. The claim follows immediately from inequality (24) in the proof of Theorem 1 in Appendix A. Note that (24) does not rely on the assumption that $\mathbb{P} > 0$. □

4. Distributionally robust predictors and prescriptors are optimal

Armed with the fundamental results of large deviations theory, we now endeavor to identify the least conservative data-driven predictors and prescriptors whose out-of-sample disappointment decays at a rate no less than some prescribed threshold $r > 0$ under any model $\mathbb{P} \in \mathcal{P}$, that is, we aim to solve the vector optimization problems (5) and (6).

4.1. Distributionally robust predictors

The relative entropy lends itself to constructing a data-driven predictor in the sense of Definition 3. We will show below that this predictor is strongly optimal in (5).

Definition 6 (Distributionally robust predictors). For any fixed threshold $r \geq 0$, we define the data-driven predictor $\hat{c}_r : X \times \mathcal{P} \to \Re$ through

$$\hat{c}_r(x, \mathbb{P}') = \max_{\mathbb{P} \in \mathcal{P}} \{c(x, \mathbb{P}) : \mathrm{I}(\mathbb{P}', \mathbb{P}) \leq r\} \qquad \forall x \in X,\ \mathbb{P}' \in \mathcal{P}. \tag{10}$$

The data-driven predictor $\hat{c}_r$ admits a distributionally robust interpretation. In fact, $\hat{c}_r(x, \mathbb{P}')$ represents the worst-case expected cost associated with the decision x, where the worst case is taken across all models $\mathbb{P} \in \mathcal{P}$ whose relative entropy distance to $\mathbb{P}'$ is at most r. Observe that the maximum in (10) is always attained because $c(x, \mathbb{P})$ is linear in $\mathbb{P}$ and the feasible set of (10) is compact, which follows from the compactness of $\mathcal{P}$ and the lower semicontinuity of the relative entropy in $\mathbb{P}$ for any fixed $\mathbb{P}'$; see Proposition 1(iii). Note also that $\hat{c}_r(x, \mathbb{P}')$ can be evaluated efficiently because (10) constitutes a convex conic optimization problem with d decision variables. Further details on how to evaluate $\hat{c}_r(x, \mathbb{P}')$ in practice are given in Appendix B.

Remark 1 (Sample average predictor). For $r = 0$ the distributionally robust predictor $\hat{c}_r$ collapses to the sample average predictor of Example 1. Indeed, because of the strict positivity of the relative entropy $\mathrm{I}(\mathbb{P}', \mathbb{P}) > 0$ for $\mathbb{P}' \neq \mathbb{P}$, see Proposition 1(i), we have that

$$\hat{c}_0(x, \mathbb{P}') = c(x, \mathbb{P}').$$

As shown in Example 2, the sample average predictor fails to offer asymptotic or finite sample guarantees of the form (3) and (4), respectively.

Remark 2 (Reverse distributionally robust predictors). The relative entropy can also be used to construct a reverse distributionally robust predictor $\check{c}_r \in \mathcal{C}$ defined through

$$\check{c}_r(x, \mathbb{P}') = \max_{\mathbb{P} \in \mathcal{P}} \{c(x, \mathbb{P}) : \mathrm{I}(\mathbb{P}, \mathbb{P}') \leq r\} \qquad \forall x \in X,\ \mathbb{P}' \in \mathcal{P}. \tag{11}$$

In contrast to $\hat{c}_r$, the reverse predictor $\check{c}_r$ fixes the *second* argument of the relative entropy and maximizes over the *first* argument. The predictors $\hat{c}_r$ and $\check{c}_r$ differ because the relative entropy fails to be symmetric. We emphasize that the reverse predictor $\check{c}_r$ has appeared often in the literature on distributionally robust optimization, see, *e.g.*, (Ben-Tal et al. 2013, Calafiore 2007, Hu and Hong 2013, Lam 2016b, Wang et al. 2016). The predictor $\hat{c}_r$ suggested here has not yet been studied extensively even though—as we will demonstrate below—it displays attractive theoretical properties that are not shared by $\check{c}_r$. The difference between $\hat{c}_r$ and $\check{c}_r$ is significant because $\mathrm{I}(\mathbb{P}', \mathbb{P})$ is finite only if $\mathbb{P}'$ is absolutely continuous with respect to $\mathbb{P}$ (*i.e.*, only if $\mathbb{P}(i) = 0$ implies $\mathbb{P}'(i) = 0$ for all $i \in \Xi$). Thus, $\check{c}_r$ hedges only against models $\mathbb{P}$ that are absolutely continuous with respect to the (observed realization of the) empirical distribution $\mathbb{P}'$. While it is clear that the empirical distribution must be absolutely continuous with respect to the data-generating distribution, however, the converse implication is generally false. Indeed, an outcome can have positive probability

even if it does not show up in a given finite time series. By taking the worst case only over models that are absolutely continuous with respect to $\mathbb{P}'$, the predictor $\check{c}_r$ potentially ignores many models that could have generated the observed data.

We first establish that $\hat{c}_r$ indeed belongs to the set $\mathcal{C}$ of all data-driven predictors, that is, the family of continuous functions mapping $X \times \mathcal{P}$ to the reals.

PROPOSITION 2 **(Continuity of $\hat{c}_r$).** *The distributionally robust predictor $\hat{c}_r : X \times \mathcal{P} \to \Re$ is continuous for any $r \geq 0$.*

Proof. Denote by $\mathcal{B}(\mathbb{P}') = \{\mathbb{P} \in \mathcal{P} : \mathrm{I}(\mathbb{P}', \mathbb{P}) \leq r\}$ the feasible set mapping of problem (10). Observe that $\mathcal{B}(\mathbb{P}')$ is non-empty for every $\mathbb{P}' \in \mathcal{P}$ by Proposition 1(i), while its graph

$$\mathrm{graph}(\mathcal{B}) = \{(\mathbb{P}', \mathbb{P}) \in \mathcal{P} \times \mathcal{P} : \mathrm{I}(\mathbb{P}', \mathbb{P}) \leq r\}$$

is convex and closed by Propositions 1(ii) and 1(iii), respectively. As $\mathcal{P}$ is compact, the closedness of $\mathrm{graph}(\mathcal{B})$ implies that the set-valued mapping $\mathcal{B}$ is upper semicontinuous (Aubin and Frankowska 1990, Proposition 1.4.8). Moreover, as $\mathcal{P}$ is a polytope, the convexity of $\mathrm{graph}(\mathcal{B})$ implies that $\mathcal{B}$ is also lower semicontinuous (Maćkowiak 2006, Theorem 3). Being both upper and lower semicontinuous, $\mathcal{B}$ constitutes in fact a continuous set-valued mapping. Moreover, $c(x, \mathbb{P})$ is jointly continuous in x and $\mathbb{P}$. Consequently, the data-driven predictor $\hat{c}_r(x, \mathbb{P}') = \max_{\mathbb{P} \in \mathcal{B}(\mathbb{P}')} c(x, \mathbb{P})$ is continuous by Berge's celebrated maximum theorem (Berge 1963, pp. 115–116). □

We now analyze the performance of the distributionally robust data-driven predictor $\hat{c}_r$ using arguments from large deviations theory. The parameter r encoding the predictor $\hat{c}_r$ captures the fundamental trade-off between out-of-sample disappointment and accuracy, which is inherent to any approach to data-driven prediction. Indeed, as r increases, the predictor $\hat{c}_r$ becomes more reliable in the sense that its out-of-sample disappointment decreases. However, increasing r also results in more conservative (pessimistically biased) predictions. In the following we will demonstrate that $\hat{c}_r$ strikes indeed an optimal balance between reliability and conservatism.

THEOREM 3 **(Feasibility of $\hat{c}_r$).** *If $r \geq 0$, then the predictor $\hat{c}_r$ is feasible in* (5).

Proof. From Proposition 2 we already know that $\hat{c}_r \in \mathcal{C}$. It remains to be shown that the out-of-sample disappointment of $\hat{c}_r$ decays at a rate of at least r. We have $c(x, \mathbb{P}) > \hat{c}_r(x, \hat{\mathbb{P}}_T)$ if and only if the estimator $\hat{\mathbb{P}}_T$ falls within the disappointment set

$$\mathcal{D}(x, \mathbb{P}) = \{\mathbb{P}' \in \mathcal{P} : c(x, \mathbb{P}) > \hat{c}_r(x, \mathbb{P}')\}.$$

Note that by the definition of $\hat{c}_r$, we have

$$\mathrm{I}(\mathbb{P}', \mathbb{P}) \leq r \quad \Longrightarrow \quad \hat{c}_r(x, \mathbb{P}') = \sup_{\mathbb{P}'' \in \mathcal{P}} \{c(x, \mathbb{P}'') : I(\mathbb{P}', \mathbb{P}'') \leq r\} \geq c(x, \mathbb{P}).$$

By contraposition, the above implication is equivalent to

$$c(x,\mathbb{P}) > \hat{c}_r(x,\mathbb{P}') \quad \Longrightarrow \quad \mathrm{I}(\mathbb{P}',\mathbb{P}) > r.$$

Therefore, $\mathcal{D}(x,\mathbb{P})$ is a subset of

$$\mathcal{I}(\mathbb{P}) = \{\mathbb{P}' \in \mathcal{P} : \mathrm{I}(\mathbb{P}',\mathbb{P}) > r\}$$

irrespective of $x \in X$. We thus have

$$\limsup_{T\to\infty} \frac{1}{T} \log \mathbb{P}^\infty \left(\hat{\mathbb{P}}_T \in \mathcal{D}(x,\mathbb{P})\right) \leq \limsup_{T\to\infty} \frac{1}{T} \log \mathbb{P}^\infty \left(\hat{\mathbb{P}}_T \in \mathcal{I}(\mathbb{P})\right) \leq -\inf_{\mathbb{P}'\in\mathcal{I}(\mathbb{P})} \mathrm{I}(\mathbb{P}',\mathbb{P}) \leq -r,$$

where the first inequality holds because $\mathcal{D}(x,\mathbb{P}) \subseteq \mathcal{I}(\mathbb{P})$, while the second inequality exploits the weak LDP upper bound (7a). Thus, $\hat{c}_r$ is feasible in (5). □

The following main theorem establishes that $\hat{c}_r$ is not only a feasible but even a strongly optimal solution for the vector optimization problem (5). This means that if an arbitrary data-driven predictor $\hat{c}$ predicts a lower expected cost than $\hat{c}_r$ even for a single estimator realization $\mathbb{P}' \in \mathcal{P}$, then $\hat{c}$ must suffer from a higher out-of-sample disappointment than $\hat{c}_r$ to first order in the exponent.

THEOREM 4 **(Optimality of $\hat{c}_r$).** *If $r > 0$, then $\hat{c}_r$ is strongly optimal in* (5).

Proof. Assume for the sake of argument that $\hat{c}_r$ fails to be a strong solution for (5). Thus, there exists a data-driven predictor $\hat{c} \in \mathcal{C}$ that is feasible in (5) but not dominated by $\hat{c}_r$, that is, $\hat{c}_r \not\preceq_{\mathcal{C}} \hat{c}$. This means that there exists $x \in X$ and $\mathbb{P}'_0 \in \mathcal{P}$ with $\hat{c}_r(x,\mathbb{P}'_0) > \hat{c}(x,\mathbb{P}'_0)$. For later reference we set $\epsilon = \hat{c}_r(x,\mathbb{P}'_0) - \hat{c}(x,\mathbb{P}'_0) > 0$. In the remainder of the proof we will demonstrate that $\hat{c}$ cannot be feasible in (5), which contradicts our initial assumption.

Let $\mathbb{P}_0 \in \mathcal{P}$ be an optimal solution of problem (10) at $\mathbb{P}' = \mathbb{P}'_0$. Thus, we have $I(\mathbb{P}'_0,\mathbb{P}_0) \leq r$ and

$$\hat{c}_r(x,\mathbb{P}'_0) = c(x,\mathbb{P}_0). \tag{12}$$

In the following we will first perturb $\mathbb{P}_0$ to obtain a model $\mathbb{P}_1$ that is $\frac{\epsilon}{2}$-suboptimal in (10) but satisfies $I(\mathbb{P}'_0,\mathbb{P}_1) < r$. Subsequently, we will perturb $\mathbb{P}_1$ to obtain a model $\mathbb{P}_2$ that is ϵ-suboptimal in (10) but satisfies $I(\mathbb{P}'_0,\mathbb{P}_2) < r$ as well as $\mathbb{P}_2 > 0$.

To construct $\mathbb{P}_1$, consider all models $\mathbb{P}(\lambda) = \lambda\mathbb{P}'_0 + (1-\lambda)\mathbb{P}_0$, $\lambda \in [0,1]$, on the line segment between $\mathbb{P}'_0$ and $\mathbb{P}_0$. As r is strictly positive, the convexity of the relative entropy implies that

$$\mathrm{I}(\mathbb{P}'_0,\mathbb{P}(\lambda)) \leq \lambda \mathrm{I}(\mathbb{P}'_0,\mathbb{P}'_0) + (1-\lambda)\mathrm{I}(\mathbb{P}'_0,\mathbb{P}_0) \leq (1-\lambda)r < r \quad \forall \lambda \in (0,1].$$

Moreover, as the expected cost $c(x,\mathbb{P}(\lambda))$ changes continuously in λ, there exists a sufficiently small $\lambda_1 \in (0,1]$ such that $\mathbb{P}_1 = \mathbb{P}(\lambda_1)$ and $r_1 = \mathrm{I}(\mathbb{P}'_0,\mathbb{P}_1)$ satisfy $0 < r_1 < r$ and

$$c(x,\mathbb{P}_0) < c(x,\mathbb{P}_1) + \frac{\epsilon}{2}.$$

To construct $\mathbb{P}_2$, we consider all models $\mathbb{P}(\lambda) = \lambda\mathbb{U} + (1-\lambda)\mathbb{P}_1,\ \lambda \in [0,1]$, on the line segment between the uniform distribution $\mathbb{U}$ on Ξ and $\mathbb{P}_1$. By the convexity of the relative entropy we have

$$\mathrm{I}(\mathbb{P}_0', \mathbb{P}(\lambda)) \leq \lambda \mathrm{I}(\mathbb{P}_0', \mathbb{U}) + (1-\lambda)\mathrm{I}(\mathbb{P}_0', \mathbb{P}_1) \leq \lambda \mathrm{I}(\mathbb{P}_0', \mathbb{U}) + (1-\lambda) r_1 \quad \forall \lambda \in [0,1].$$

As $r_1 < r$ and the expected cost $c(x, \mathbb{P}(\lambda))$ changes continuously in λ, there exists a sufficiently small $\lambda_2 \in (0,1]$ such that $\mathbb{P}_2 = \mathbb{P}(\lambda_2)$ and $r_2 = \mathrm{I}(\mathbb{P}_0', \mathbb{P}_2)$ satisfy $0 < r_2 < r$, $\mathbb{P}_2 > 0$ and

$$c(x, \mathbb{P}_0) < c(x, \mathbb{P}_2) + \epsilon. \tag{13}$$

In summary, we thus have

$$\hat{c}(x, \mathbb{P}_0') = \hat{c}_r(x, \mathbb{P}_0') - \epsilon = c(x, \mathbb{P}_0) - \epsilon < c(x, \mathbb{P}_2) \leq \hat{c}_r(x, \mathbb{P}_0'), \tag{14}$$

where the first equality follows from the definition of ϵ, and the second equality exploits (12). Moreover, the strict inequality holds due to (13), and the weak inequality follows from the definition of $\hat{c}_r$ and the fact that $\mathrm{I}(\mathbb{P}_0', \mathbb{P}_2) = r_2 < r$.

In the remainder of the proof we will argue that the prediction disappointment $\mathbb{P}_2^\infty(c(x, \mathbb{P}_2) > \hat{c}(x, \hat{\mathbb{P}}_T))$ under model $\mathbb{P}_2$ decays at a rate of at most $r_2 < r$ as the sample size T tends to infinity. In analogy to the proof of Theorem 3, we define the set of disappointing estimator realizations as

$$\mathcal{D}(x, \mathbb{P}_2) = \left\{ \mathbb{P}' \in \mathcal{P} \ : \ c(x, \mathbb{P}_2) > \hat{c}(x, \mathbb{P}') \right\}.$$

This set contains $\mathbb{P}_0'$ due to the strict inequality in (14). Moreover, as $\hat{c} \in \mathcal{C}$ is continuous, $\mathcal{D}(x, \mathbb{P}_2)$ is an open and, *a fortiori*, I-continuous subset of $\mathcal{P}$. Thus, we find

$$\inf_{\mathbb{P}' \in \operatorname{int} \mathcal{D}(x, \mathbb{P}_2)} \mathrm{I}(\mathbb{P}', \mathbb{P}_2) = \inf_{\mathbb{P}' \in \mathcal{D}(x, \mathbb{P}_2)} \mathrm{I}(\mathbb{P}', \mathbb{P}_2) \leq \mathrm{I}(\mathbb{P}_0', \mathbb{P}_2) = r_2,$$

where the inequality holds because $\mathbb{P}_0' \in \mathcal{D}(x, \mathbb{P}_0)$, and the last equality follows from the definition of r_2. As the empirical distributions $\{\hat{\mathbb{P}}_T\}_{T \in \mathbb{N}}$ obey the LDP lower bound (7b) under $\mathbb{P}_2 > 0$, we finally conclude that

$$-r < -r_2 \leq -\inf_{\mathbb{P}' \in \operatorname{int} \mathcal{D}(x, \mathbb{P}_2)} \mathrm{I}(\mathbb{P}', \mathbb{P}_2) \leq \liminf_{T \to \infty} \frac{1}{T} \log \mathbb{P}_2^\infty \left(\hat{\mathbb{P}}_T \in \mathcal{D}(x, \mathbb{P}_2) \right).$$

The above chain of inequalities implies, however, that $\hat{c}$ is infeasible in problem (5). This contradicts our initial assumption, and thus, $\hat{c}_r$ must indeed be a strong solution of (5). □

Theorem 4 asserts that the distributionally robust predictor $\hat{c}_r$ is optimal among all data-driven predictors representable as continuous functions of the empirical distribution $\hat{\mathbb{P}}_T$. That is, any attempt to make it less conservative invariable increases the out-of-sample prediction disappointment. We remark that the class of predictors which depend on the data only through $\hat{\mathbb{P}}_T$ is vast.

These predictors constitute arbitrary continuous functions of the data that are independent of the order in which the samples were observed. As the samples are independent and identically distributed, there are in fact no meaningful data-driven predictors that display a more general dependence on the data.

Note that in the above discussion all guarantees are fundamentally asymptotic in nature. Using Theorem 2 one can show, however, that $\hat{c}_r$ also satisfies finite sample guarantees.

THEOREM 5 **(Finite sample guarantee)**. *The out-of-sample disappointment of the distributionally robust predictor $\hat{c}_r$ enjoys the following finite sample guarantee under any model $\mathbb{P} \in \mathcal{P}$ and for any $x \in X$.*

$$\mathbb{P}^\infty \left(c(x, \mathbb{P}) > \hat{c}_r(x, \hat{\mathbb{P}}_T) \right) \leq (T+1)^d e^{-rT} \quad \forall T \in \mathbb{N} \tag{15}$$

Proof. The proof of this result widely parallels that of Theorem 3 but uses the strong LDP upper bound (9) in lieu of the weak upper bound (7a). Details are omitted for brevity. □

4.2. Distributionally robust prescriptors

The distributionally robust predictor $\hat{c}_r$ of Definition 6 induces a corresponding prescriptor.

DEFINITION 7 (DISTRIBUTIONALLY ROBUST PRESCRIPTORS). Denote by $\hat{c}_r$, $r \geq 0$, the distributionally robust data-driven predictor of Definition 6. We can then define the data-driven prescriptor $\hat{x}_r : \mathcal{P} \to X$ as a quasi-continuous function satisfying

$$\hat{x}_r(\mathbb{P}') \in \arg\min_{x \in X} \hat{c}_r(x, \mathbb{P}') \quad \forall \mathbb{P}' \in \mathcal{P}. \tag{16}$$

Note that the minimum in (16) is attained because X is compact and $\hat{c}_r$ is continuous due to Proposition 2. Thus, there exists at least one function $\hat{x}_r$ satisfying (16). In the next proposition we argue that this function can be chosen to be quasi-continuous as desired.

PROPOSITION 3 **(Quasi-continuity of $\hat{x}_r$)**. *If $r \geq 0$, then there exists a quasi-continuous data-driven predictor $\hat{x}_r$ satisfying* (16).

Proof. Denote by $\Gamma(\mathbb{P}') = \arg\min_{x \in X} \hat{c}_r(x, \mathbb{P}')$ the argmin-mapping of problem (10), and observe that $\Gamma(\mathbb{P}')$ is compact and non-empty for every $\mathbb{P}' \in \mathcal{P}$ because $\hat{c}_r$ is continuous and X is compact. As X is independent of $\mathbb{P}'$, Berge's maximum theorem (Berge 1963, pp. 115–116) further implies that Γ is upper semicontinuous, while (Matejdes 1987, Theorem 5.3) guarantees that there exists a quasi-continuous function $\hat{x}_r : \mathcal{P} \to X$ with $\hat{x}_r(\mathbb{P}') \in \Gamma(\mathbb{P}')$ for all $\mathbb{P}' \in \mathcal{P}$. □

Propositions 2 and 3 imply that $(\hat{c}_r, \hat{x}_r)$ belongs to the family $\mathcal{X}$ of all data-driven predictor-prescriptor-pairs. Using a similar reasoning as in Theorem 3, we now demonstrate that the out-of-sample disappointment of $\hat{x}_r$ decays at rate at least r as T tends to infinity. Thus, $\hat{x}_r$ provides trustworthy prescriptions.

THEOREM 6 **(Feasibility of $(\hat{c}_r, \hat{x}_r)$).** *If $r \geq 0$, then the predictor-prescriptor pair $(\hat{c}_r, \hat{x}_r)$ is feasible in* (6).

Proof. Propositions 2 and 3 imply that $(\hat{c}_r, \hat{x}_r) \in \mathcal{X}$. It remains to be shown that the out-of-sample disappointment of $\hat{x}_r$ decays at a rate of at least r. To this end, define $\mathcal{D}(x, \mathbb{P})$ and $\mathcal{I}(\mathbb{P})$ as in the proof of Theorem 3, and recall that $\mathcal{D}(x, \mathbb{P}) \subseteq \mathcal{I}(\mathbb{P})$ for every decision $x \in X$ and model $\mathbb{P} \in \mathcal{P}$. Thus, for every fixed estimator realization $\mathbb{P}' \in \mathcal{P}$ the following implication holds

$$\begin{array}{lll} c(\hat{x}_r(\mathbb{P}'), \mathbb{P}) > \hat{c}_r(\hat{x}_r(\mathbb{P}'), \mathbb{P}') & \Longrightarrow & \exists x \in X \text{ with } c(x, \mathbb{P}) > \hat{c}_r(x, \mathbb{P}') \\ & \Longrightarrow & \mathbb{P}' \in \cup_{x \in X} \mathcal{D}(x, \mathbb{P}) \\ & \Longrightarrow & \mathbb{P}' \in \mathcal{I}(\mathbb{P}), \end{array}$$

which in turn implies

$$\limsup_{T \to \infty} \frac{1}{T} \log \mathbb{P}^\infty \left(c(\hat{x}_r(\hat{\mathbb{P}}_T), \mathbb{P}) > \hat{c}_r(\hat{x}_r(\hat{\mathbb{P}}_T), \hat{\mathbb{P}}_T) \right) \leq \limsup_{T \to \infty} \frac{1}{T} \log \mathbb{P}^\infty \left(\hat{\mathbb{P}}_T \in \mathcal{I}(\mathbb{P}) \right) \leq -r$$

for every model $\mathbb{P} \in \mathcal{P}$. Note that the second inequality in the above expression has already been established in the proof of Theorem 3. Thus, the claim follows. □

Next, we argue that $(\hat{c}_r, \hat{x}_r)$ is a strongly optimal solution for the vector optimization problem (6).

THEOREM 7 **(Optimality of $(\hat{c}_r, \hat{x}_r)$).** *If $r > 0$, then $(\hat{c}_r, \hat{x}_r)$ is strongly optimal in* (6).

Proof. Assume for the sake of argument that $(\hat{c}_r, \hat{x}_r)$ fails to be a strong solution for (6). Thus, there exists a data-driven prescriptor $(\hat{c}, \hat{x}) \in \mathcal{X}$ that is feasible in (6) but not dominated by $(\hat{c}_r, \hat{x}_r)$, that is, $(\hat{c}_r, \hat{x}_r) \not\preceq_{\mathcal{X}} (\hat{c}, \hat{x})$. This means that there exists $\mathbb{P}'_0 \in \mathcal{P}$ with $\hat{c}_r(\hat{x}_r(\mathbb{P}'_0), \mathbb{P}'_0) > \hat{c}(\hat{x}(\mathbb{P}'_0), \mathbb{P}'_0)$. As X is compact and $\hat{c}$ is continuous, the cost $\hat{c}(\hat{x}(\mathbb{P}'), \mathbb{P}')$ of the prescriptor $\hat{x}$ under the corresponding predictor $\hat{c}$ is continuous in $\mathbb{P}'$ (Rockafellar and Wets 1998, Theorem 1.17(c)). Similarly, $\hat{c}_r(\hat{x}_r(\mathbb{P}'), \mathbb{P}')$ is continuous in $\mathbb{P}'$. Recall also that $\hat{x}$ is quasi-continuous and therefore continuous on a dense subset of $\mathcal{P}$ (Bledsoe 1952). Thus, we may assume without loss of generality that $\hat{x}$ is continuous at $\mathbb{P}'_0$. For later reference we set $\epsilon = \hat{c}_r(\hat{x}(\mathbb{P}'_0), \mathbb{P}'_0) - \hat{c}(\hat{x}(\mathbb{P}'_0), \mathbb{P}'_0) > 0$.

In the remainder of the proof we will demonstrate that $(\hat{c}, \hat{x})$ cannot be feasible in (6), which contradicts our initial assumption. To this end, let $\mathbb{P}_0 \in \mathcal{P}$ be an optimal solution of problem (10) at $x = \hat{x}(\mathbb{P}'_0)$ and $\mathbb{P}' = \mathbb{P}'_0$. Thus, we have $I(\mathbb{P}'_0, \mathbb{P}_0) \leq r$ and

$$\hat{c}_r(\hat{x}(\mathbb{P}'_0), \mathbb{P}'_0) = c(\hat{x}(\mathbb{P}'_0), \mathbb{P}_0). \tag{17}$$

Next, we first perturb $\mathbb{P}_0$ to obtain a model $\mathbb{P}_1$ that is strictly $\frac{\epsilon}{2}$-suboptimal in (10) but satisfies $I(\mathbb{P}'_0, \mathbb{P}_1) = r_1 < r$. Subsequently, we perturb $\mathbb{P}_1$ to obtain a model $\mathbb{P}_2$ that is strictly ϵ-suboptimal

in (10) but satisfies $I(\mathbb{P}_0', \mathbb{P}_2) = r_2 < r$ as well and $\mathbb{P}_2 > 0$. The distributions $\mathbb{P}_1$ and $\mathbb{P}_2$ can be constructed exactly as in the proof of Theorem 4. Details are omitted for brevity. Thus, we find

$$\hat{c}(\hat{x}(\mathbb{P}_0'), \mathbb{P}_0') = \hat{c}_r(\hat{x}(\mathbb{P}_0'), \mathbb{P}_0') - \epsilon = c(\hat{x}(\mathbb{P}_0'), \mathbb{P}_0) - \epsilon < c(\hat{x}(\mathbb{P}_0'), \mathbb{P}_2) \leq \hat{c}_r(\hat{x}(\mathbb{P}_0'), \mathbb{P}_0'), \tag{18}$$

where the first equality follows from the definition of ϵ, and the second equality exploits (17). Moreover, the strict inequality holds because $\mathbb{P}_2$ is strictly ϵ-suboptimal in (10), while the weak inequality follows from the definition of $\hat{c}_r$ and the fact that $I(\mathbb{P}_0', \mathbb{P}_2) = r_2 < r$.

It remains to be shown that the prediction disappointment $\mathbb{P}_2^\infty(c(\hat{x}(\hat{\mathbb{P}}_T), \mathbb{P}_2) > \hat{c}(\hat{x}(\hat{\mathbb{P}}_T), \hat{\mathbb{P}}_T))$ under model $\mathbb{P}_2$ decays at a rate of at most $r_2 < r$ as the sample size T tends to infinity. To this end, we define the set of disappointing estimator realizations as

$$\mathcal{D}(\mathbb{P}_2) = \{\mathbb{P}' \in \mathcal{P} \ : \ c(\hat{x}(\mathbb{P}'), \mathbb{P}_2) > \hat{c}(\hat{x}(\mathbb{P}'), \mathbb{P}')\}.$$

This set contains $\mathbb{P}_0'$ due to the strict inequality in (18). Recall now that $\hat{x}$ is continuous at $\mathbb{P}' = \mathbb{P}_0'$ due to our choice of $\mathbb{P}_0'$. As the predictors $\hat{c}$ and $\hat{c}_r$ are both continuous on their entire domain, the compositions $\hat{c}(\hat{x}(\mathbb{P}'), \mathbb{P}')$ and $c(\hat{x}(\mathbb{P}'), \mathbb{P}_2)$ are both continuous at $\mathbb{P}' = \mathbb{P}_0'$. This implies that $\mathbb{P}_0'$ belongs actually to the interior of $\mathcal{D}(\mathbb{P}_2)$. Thus, we find

$$\inf_{\mathbb{P}' \in \operatorname{int} \mathcal{D}(\mathbb{P}_2)} \mathrm{I}(\mathbb{P}', \mathbb{P}_2) \leq \mathrm{I}(\mathbb{P}_0', \mathbb{P}_2) = r_2,$$

where the last equality follows from the definition of r_2. As the empirical distributions $\{\hat{\mathbb{P}}_T\}_{T \in \mathbb{N}}$ obey the LDP lower bound (7b) under $\mathbb{P}_2 > 0$, we finally conclude that

$$-r < -r_2 \leq -\inf_{\mathbb{P}' \in \operatorname{int} \mathcal{D}(\mathbb{P}_2)} \mathrm{I}(\mathbb{P}', \mathbb{P}_2) \leq \liminf_{T \to \infty} \frac{1}{T} \log \mathbb{P}_2^\infty \left(\hat{\mathbb{P}}_T \in \mathcal{D}(\mathbb{P}_2) \right).$$

The above chain of inequalities implies, however, that $(\hat{c}, \hat{x})$ is infeasible in problem (6). This contradicts our initial assumption, and thus, $(\hat{c}_r, \hat{x}_r)$ must indeed be a strong solution of (6). □

All guarantees discussed so far are asymptotic in nature. As in the case of the predictor $\hat{c}_r$, however, the prescriptor $\hat{x}_r$ can also be shown to satisfy finite sample guarantees.

THEOREM 8 **(Finite sample guarantee)**. *The out-of-sample disappointment of the distributionally robust prescriptor $\hat{x}_r$ enjoys the following finite sample guarantee under any model $\mathbb{P} \in \mathcal{P}$.*

$$\mathbb{P}^\infty \left(c(\hat{x}_r(\hat{\mathbb{P}}_T), \mathbb{P}) > \hat{c}_r(\hat{x}_r(\hat{\mathbb{P}}_T), \hat{\mathbb{P}}_T) \right) \leq (T+1)^d e^{-rT} \quad \forall T \in \mathbb{N} \tag{19}$$

Proof. The proof of this result parallels those of Theorems 3 and 6 but uses the strong LDP upper bound (9) in lieu of the weak upper bound (7a). Details are omitted for brevity. □

We stress that the finite sample guarantees of Theorems 5 and 8 as well as the strong optimality properties portrayed in Theorems 4 and 7 are independent of a particular dataset. They guarantee that $\hat{c}_r$ and $\hat{x}_r$ provide trustworthy predictions and prescriptions, respectively, *before* the data is revealed. In Section 5 we investigate what can be said about the quality of these data-driven solutions *after* the data is revealed.

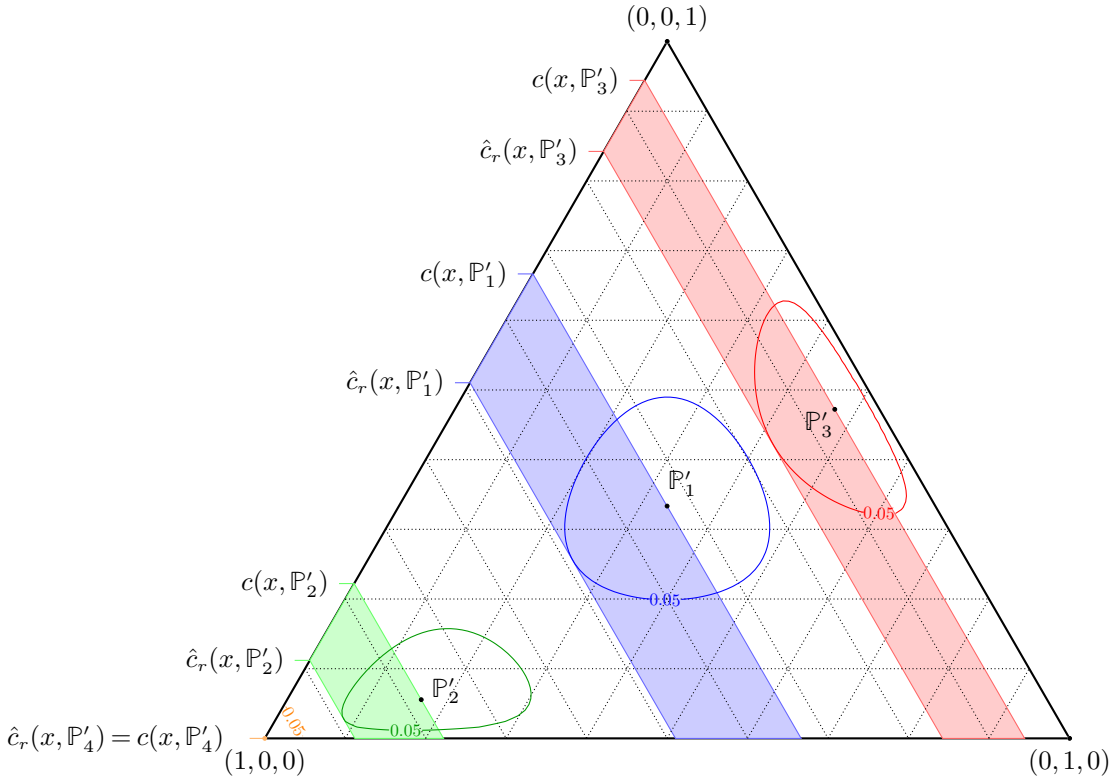

Figure 2 **Visualization of the pseudo-norm ball $\mathcal{B}(\mathbb{P}')=\{\mathbb{P}\in\mathcal{P} \,:\, \mathrm{I}(\mathbb{P}',\mathbb{P})\leq r\}$ around different centers $\mathbb{P}'$. The conservatism of $\hat{c}_r$ at a particular center is captured by the width of the corresponding colored stripe.**

5. The price of combating out-of-sample disappointment

The rate at which the out-of-sample disappointment of a data-driven predictor $\hat{c}$ decays is an *a priori* property of $\hat{c}$ as a function mapping observations to predictions. It depends on the unknown data-generating distribution but is independent of the observed dataset. Thus, it can in principle be determined before observing a single data point. Note that the vector optimization problem (5) focuses exclusively on predictors $\hat{c}$ whose out-of-sample disappointment is guaranteed to decay at a rate of at least r under every possible data-generating distribution.

The conservatism $\hat{c}(x,\mathbb{P}')-c(x,\mathbb{P}')$ of the data-driven predictor $\hat{c}$ at $x\in X$ and $\mathbb{P}'\in\mathcal{P}$ constitutes an *a posteriori* property of $\hat{c}$ and quantifies the difference between the predicted cost of x and the cost of x under the predicted model $\mathbb{P}'$. The conservatism can only be determined after a particular realization $i_t\in\Xi$ of the random variable ξ_t has been observed for each $t=1,\ldots,T$, which gives rise to a particular realization $\mathbb{P}'$ of the estimator $\hat{\mathbb{P}}_T$. Note that the vector optimization problem (5) minimizes the conservatism of $\hat{c}$ at all $x\in X$ and $\mathbb{P}'\in\mathcal{P}$.

We can view the conservatism of a data-driven predictor as the price paid for the guarantee that its out-of-sample disappointment fades sufficiently quickly. In the remainder of this section we will investigate the fundamental trade-off between the *a priori* and *a posteriori* properties of

data-driven predictors in the special case of $\hat{c}_r$. Recall from Defintion 6 that $\hat{c}_r$ represents the worst-case expected cost over all models in the pseudo-norm ball $\mathcal{B}(\mathbb{P}') = \{\mathbb{P} \in \mathcal{P} \,:\, \mathrm{I}(\mathbb{P}', \mathbb{P}) \leq r\}$ centered at $\mathbb{P}'$. The size and shape of $\mathcal{B}(\mathbb{P}')$ determine the conservatism of $\hat{c}_r$ and thus the price to be paid for the guarantee that the out-of-sample disappointment of $\hat{c}_r$ decays at rate r. Like $\mathcal{B}(\mathbb{P}')$, the conservatism thus depends both on the desired decay rate r and on the observed dataset. This suggests that not all data-sets of the same cardinality are equally valuable. The smaller the conservatism $\hat{c}_r(x, \mathbb{P}') - c(x, \mathbb{P}')$ needed for the decay rate r, the more valuable is the observation $\mathbb{P}'$. Figure 2 shows four pseudo-norm balls $\mathcal{B}(\mathbb{P}')$ with radius $r = 5\%$ around four distinct empirical distributions $\mathbb{P}'_1 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, $\mathbb{P}'_2 = (\frac{14}{18}, \frac{3}{18}, \frac{1}{18})$, $\mathbb{P}'_3 = (\frac{2}{36}, \frac{17}{36}, \frac{17}{36})$ and $\mathbb{P}'_4 = (1, 0, 0)$ in a ternary diagram. The colored stripes visualize the conservatism of the distributionally robust predictor $\hat{c}_r$ at the different empirical distributions corresponding to the cost function $\gamma(x, \xi) = \mathbb{1}_{\xi_1 = 1}$. Observe that the conservatism can change dramatically with $\mathbb{P}'$. In the extreme case when $\mathbb{P}' = \mathbb{P}'_4$ the pseudo-norm ball reduces to the subsimplex $\mathcal{B}(\mathbb{P}'_4) = \{\mathbb{P} \in \mathcal{P} \,:\, e^{-r} \leq \mathbb{P}(1) \leq 1\}$, in which case the conservatism drops to zero. In view of this discussion, a dataset with empirical distribution $\mathbb{P}'_3$ is more valuable than one with empirical distribution $\mathbb{P}'_1$. The most valuable datasets are those with empirical distribution $\mathbb{P}'_4$.

To gain deeper insights into the dependence of the conservatism on r and $\mathbb{P}'$, we now investigate the asymptotics of the pseudo-norm ball $\mathcal{B}(\mathbb{P}')$ as r drops to 0. We start by observing that the relative entropy is closely related to the observed Fischer information and the Pearson divergence.

Definition 8 (Observed Fisher information). The observed Fisher information corresponding to the sequence of observations $\{i_t\}_{t=1}^T$ and model $\mathbb{P} \in \mathcal{P}$, $\mathbb{P} > 0$, is defined as the matrix

$$F(\mathbb{P}) = -\frac{1}{T} \nabla^2_{\mathbb{P}} \left[\log \mathbb{P}^\infty (\xi_1 = i_1, \ldots, \xi_T = i_T) \right] \in \mathbb{S}^d.$$

Definition 9 (Pearson divergence). The Pearson divergence (also known as χ^2-distance) of an estimator realization $\mathbb{P}' \in \mathcal{P}$, $\mathbb{P}' > 0$, with respect to a model $\mathbb{P} \in \mathcal{P}$ is defined as

$$\chi^2(\mathbb{P}', \mathbb{P}) = \sum_{i \in \Xi} \frac{(\mathbb{P}'(i) - \mathbb{P}(i))^2}{\mathbb{P}'(i)}.$$

The following proposition shows that for $\mathbb{P}'$ close to $\mathbb{P}$ the relative entropy reduces to a quadratic form whose Hessian matrix coincides with the observed Fisher information. This elementary result is described in Cover and Thomas (2006, Chapter 11). We include a short proof in Appendix A to keep the paper self-contained.

Proposition 4 **(Taylor expansion of the relative entropy)**. *If $\mathbb{P}'$ is the empirical distribution corresponding to the observations $\{i_t\}_{t=1}^T$ and $\mathbb{P}' > 0$, then the relative entropy satisfies*

$$\mathrm{I}(\mathbb{P}', \mathbb{P}) = \frac{1}{2} (\mathbb{P} - \mathbb{P}')^\top F(\mathbb{P}') (\mathbb{P} - \mathbb{P}') + o(\|\mathbb{P} - \mathbb{P}'\|^2) = \frac{1}{2} \chi^2(\mathbb{P}', \mathbb{P}) + o(\|\mathbb{P} - \mathbb{P}'\|^2), \tag{20}$$

where the observed Fisher information under model $\mathbb{P} = \mathbb{P}'$ is given by $F(\mathbb{P}') = \mathrm{diag}(\mathbb{P}')^{-1} \succ 0$.

Proposition 4 implies that for small decay rates r the relative entropy ball $\mathcal{B}(\mathbb{P}')$ can be approximated by the Pearson divergence ellipsoid $\mathcal{B}_{\chi^2}(\mathbb{P}') = \left\{\mathbb{P} \in \mathcal{P} \ : \ \frac{1}{2}\chi^2(\mathbb{P}', \mathbb{P}) \leq r\right\}$ whose shape matrix is given by the inverse of the observed Fischer information matrix. This insight further suggests that for small r the predictor $\hat{c}_r$ is approximately equal to a distributionally robust predictor that replaces the relative entropy ball with the Pearson divergence ellipsoid.

THEOREM 9 **(Rate expansion)**. *Let $\mathbb{P}'$ be the empirical distribution corresponding to the observations $\{i_t\}_{t=1}^T$, and assume that $\mathbb{P}' > 0$. Then, for small decay rates r we have*

$$\hat{c}_r(x, \mathbb{P}') = \max_{\mathbb{P} \in \mathcal{P}} \left\{c(x, \mathbb{P}) \ : \ \chi^2(\mathbb{P}', \mathbb{P}) \leq 2r\right\} + o(\sqrt{r}) \tag{21a}$$

$$= c(x, \mathbb{P}') + \sqrt{\sum_{i \in \Xi} \mathbb{P}'(i) \left(\gamma(x, j) - \sum_{j \in \Xi} \mathbb{P}'(j)\gamma(x, j)\right)^2} \cdot \sqrt{2r} + o(\sqrt{r}). \tag{21b}$$

Proof. Theorem 9 is a corollary of a more general result by Lam (2016a, Theorem 7). We present an elementary proof in Appendix A for completeness. □

Theorem 9 has several merits. On the one hand, it asserts that the conservatism of the optimal predictor $\hat{c}_r$ grows with the square root of r. The marginal price of combatting the out-of-sample disappointment is therefore decreasing in r. On the other hand, Theorem 9 reveals that the conservatism of $\hat{c}_r$ is proportional to the empricical standard deviation of $\gamma(x, \xi)$. That is, datasets under which the empirical standard deviation of $\gamma(x, \xi)$ is lower appear more valuable.

The distributionally robust predictor $\hat{c}_r$ is of interest because it is optimal in the vector optimization problem (5). The optimality of $\hat{c}_r$ offers via Theorem 9 a possible explanation for the popularity of the Pearson divergence predictor in (21a), which has attracted considerable attention in the recent literature, see, *e.g.*, (Bertsimas et al. 2014, Bayraksan and Love 2015, Jiang and Guan 2016b, Postek et al. 2016). Theorem 9 also offers an explanation for the popularity of the Markowitz-type predictor in (21b). We stress that even though the Pearson divergence and Markowitz-type predictors are asymptotically equal to $\hat{c}_r$, they are *not* optimal in (5). Also, when r is large or not all outcomes $i \in \Xi$ have been observed, then the Pearson divergence and Markowitz-type predictors can differ substantially from $\hat{c}_r$. However, they may be preferred to $\hat{c}_r$ for reasons of computational tractability.

Acknowledgments The first and second authors are generously supported by the Early Postdoc.Mobility fellowships P2EZP2_165226 and P2EZP2_165264 of the Swiss National Science Foundation, respectively. The third author gratefully acknowledges funding from the Swiss National Science Foundation under grant BSCGI0_157733.

References

J.-P. Aubin and H. Frankowska. *Set-Valued Analysis*. Springer, 1990.

G. Bayraksan and D. K. Love. Data-driven stochastic programming using phi-divergences. In *The Operations Research Revolution*, chapter 1, pages 1–19. INFORMS, 2015.

A. Ben-Tal, D. Den Hertog, A. De Waegenaere, B. Melenberg, and G. Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.

C. Berge. *Topological Spaces: Including a Treatment of Multi-Valued Functions, Vector Spaces, and Convexity*. Courier Corporation, 1963.

D. Bertsimas, V. Gupta, and N. Kallus. Robust SAA. *arXiv preprint arXiv:1408.4445*, 2014.

W. Bledsoe. Neighborly functions. *Proceedings of the American Mathematical Society*, 3:114–115, 1952.

G.C. Calafiore. Ambiguous risk measures and optimal robust portfolios. *SIAM Journal on Optimization*, 18 (3):853–877, 2007.

V. Chandrasekaran and P. Shah. Relative entropy optimization and its applications. *Mathematical Programming*, 161(1-2):1–32, 2017.

P.R. Chares. *Cones and Interior-Point Algorithms for Structured Convex Optimization Involving Powers and Exponentials*. PhD thesis, Université Catholique de Louvain, 2009.

T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley, 2006.

E. Delage and Y. Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612, 2010.

A. Domahidi, E. Chu, and S. Boyd. ECOS: An SOCP solver for embedded systems. In *European Control Conference (ECC)*, pages 3071–3076, 2013.

J. Dupačová and R. Wets. Asymptotic behavior of statistical estimators and of optimal solutions of stochastic optimization problems. *The Annals of Statistics*, 16(4):1517–1549, 1988.

E. Erdoğan and G. Iyengar. Ambiguous chance constrained problems and robust optimization. *Mathematical Programming*, 107(1–2):37–61, 2006.

V. Gupta. Near-optimal ambiguity sets for distributionally robust optimization. *Available from Optimization Online*, 2015.

Z. Hu and L.J. Hong. Kullback-Leibler divergence constrained distributionally robust optimization. *Available from Optimization Online*, 2013.

R. Jiang and Y. Guan. Risk-averse two-stage stochastic program with distributional ambiguity. *Available from Optimization Online*, 2016a.

R. Jiang and Y. Guan. Data-driven chance constrained stochastic program. *Mathematical Programming*, 158(1-2):291–327, 2016b.

S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22 (1):79–86, 1951.

H. Lam. Recovering best statistical guarantees via the empirical divergence-based distributionally robust optimization. *arXiv preprint arXiv:1605.09349*, 2016a.

H. Lam. Robust sensitivity analysis for stochastic systems. *Mathematics of Operations Research*, 41(4): 1248–1275, 2016b.

O.P. Le Maître and O.M. Knio. *Introduction: Uncertainty Quantification and Propagation*. Springer, 2010.

P. Maćkowiak. Some remarks on lower hemicontinuity of convex multivalued mappings. *Economic Theory*, 28(1):227–233, 2006.

M. Matejdes. Sur les sélecteurs des multifonctions. *Mathematica Slovaca*, 37(1):111–124, 1987.

R. O. Michaud. The Markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal*, 45(1):31–42, 1989.

P. Mohajerin Esfahani and D. Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *arXiv preprint arXiv:1505.05116*, 2015.

Y. Nesterov and A. Nemirovskii. *Interior-Point Polynomial Algorithms in Convex Programming*. SIAM, 1994.

A.B. Owen. Empirical likelihood ratio confidence intervals for a single functional. *Biometrika*, 75(2):237–249, 1988.

P. Parpas, B. Ustun, M. Webster, and Q. Tran. Importance sampling in stochastic programming: A Markov Chain Monte Carlo approach. *INFORMS Journal on Computing*, 27(2):358–377, 2015.

G. Pflug and D. Wozabal. Ambiguity in portfolio selection. *Quantitative Finance*, 7(4):435–442, 2007.

K. Postek, D. den Hertog, and B. Melenberg. Computationally tractable counterparts of distributionally robust constraints on risk measures. *SIAM Review*, 58(4):603–650, 2016.

R.T. Rockafellar and R.J.-B. Wets. *Variational Analysis*. Springer, 1998.

S.A. Serrano. *Algorithms for Unsymmetric Cone Optimization and an Implementation for Problems with the Exponential Cone*. PhD thesis, Stanford University, 2015.

A. Shapiro, D. Dentcheva, and A. Ruszczyńsk. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, 2014.

J.E. Smith and R.L. Winkler. The optimizer's curse: Skepticism and postdecision surprise in decision analysis. *Management Science*, 52(3):311–322, 2006.

H. Sun and H. Xu. Convergence analysis for distributionally robust optimization and equilibrium problems. *Mathematics of Operations Research*, 41(2):377–401, 2016.

Z. Wang, P.W. Glynn, and Y. Ye. Likelihood robust optimization for data-driven newsvendor problems. *Computational Management Science*, 12(2):241–261, 2016.

C. Zhao and Y. Guan. Data-driven risk-averse stochastic optimization with Wasserstein metric. *Available from Optimization Online*, 2015.

Appendix A: Proofs

Proof of Theorem 1. Let $i_t \in \Xi$ be a particular realization of the random variable ξ_t for each $t = 1, \ldots, T$, and denote by $\mathbb{P}'$ the realization of the estimator $\hat{\mathbb{P}}_T$ corresponding to the sequence $\{i_t\}_{t=1}^T$. The probability of observing this sequence (in the given order) under model $\mathbb{P}$ can be expressed in terms of $\mathbb{P}'$ as

$$\mathbb{P}^\infty(\xi_1 = i_1, \ldots, \xi_T = i_T) = \prod_{i \in \Xi} \mathbb{P}(i)^{T\mathbb{P}'(i)} = e^{T \sum_{i \in \Xi} \mathbb{P}'(i) \log \mathbb{P}(i)}. \tag{22}$$

Set $\mathcal{P}_T = \mathcal{P} \cap \{0, 1/T, \ldots, (T-1)/T, 1\}^d$ and note that $\mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{P}_T) = 1$. By construction, the cardinality of $\mathcal{P}_T$ is bounded above by $(T+1)^d$.

In the following, we denote the set of all sample paths in Ξ^T that give rise to the same empirical distribution $\mathbb{P}' \in \mathcal{P}_T$ by $C_T(\mathbb{P}')$. The cardinality of $C_T(\mathbb{P}')$ coincides with the number of sample paths that visit state i exactly $T \cdot \mathbb{P}'(i)$ times for all $i \in \Xi$, that is, we have

$$|C_T(\mathbb{P}')| = \frac{T!}{\prod_{i \in \Xi} (T \cdot \mathbb{P}'(i))!}.$$

Stirling's approximation for factorials allows us to bound the cardinality of $C_T(\mathbb{P}')$ from both sides in terms of the entropy $H(\mathbb{P}') = -\sum_{i=1}^d \mathbb{P}'(i) \log \mathbb{P}'(i)$ of the empirical distribution $\mathbb{P}'$, that is,

$$(T+1)^{-d} e^{T H(\mathbb{P}')} \leq |C_T(\mathbb{P}')| \leq e^{T H(\mathbb{P}')}. \tag{23}$$

An elementary proof of these inequalities that does not involve Stirling's approximation is given by Cover and Thomas (2006, Theorem 12.1.3).

Select an arbitrary Borel set $\mathcal{D} \subseteq \mathcal{P}$. For any $T \in \mathbb{N}$, we thus have

$$\begin{aligned}
\mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{D}) &= \sum_{\mathbb{P}' \in \mathcal{D} \cap \mathcal{P}_T} \mathbb{P}^\infty(\hat{\mathbb{P}}_T = \mathbb{P}') \\
&\leq (T+1)^d \cdot \max_{\mathbb{P}' \in \mathcal{D} \cap \mathcal{P}_T} \mathbb{P}^\infty(\hat{\mathbb{P}}_T = \mathbb{P}') \\
&\leq (T+1)^d \cdot \max_{\mathbb{P}' \in \mathcal{D} \cap \mathcal{P}_T} |C_T(\mathbb{P}')|\, e^{T \sum_{i \in \Xi} \mathbb{P}'(i) \log \mathbb{P}(i)} \\
&\leq (T+1)^d \cdot e^{-T \min_{\mathbb{P}' \in \mathcal{D} \cap \mathcal{P}_T} \mathrm{I}(\mathbb{P}', \mathbb{P})} \\
&\leq (T+1)^d \cdot e^{-T \inf_{\mathbb{P}' \in \mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P})},
\end{aligned}$$

where the first inequality exploits the estimate $|\mathcal{P}_T| \leq (T+1)^d$, the second inequality holds due to (22) and the definition of $C_T(\mathbb{P}')$, and the third inequality follows from the upper estimate in (23). Taking logarithms on both sides of the above expression and dividing by T yields

$$\frac{1}{T} \log \mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{D}) \leq \frac{d \log(T+1)}{T} - \inf_{\mathbb{P}' \in \mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P}). \tag{24}$$

Note that the finite sample bound (24) does not rely on any properties of the set $\mathcal{D}$ besides measurability. The asymptotic upper bound (7a) is obtained by taking the limit superior as T tends to infinity on both sides of (24).

As for the lower bound (7b), recall that $\mathrm{I}(\mathbb{P}', \mathbb{P})$ is continuous in $\mathbb{P}'$ as $\mathbb{P} > 0$, see Proposition 1(iii), and note that $\bigcup_{T\in\mathbb{N}} \mathcal{P}_T$ is dense in $\operatorname{int}\mathcal{D}$. Thus, there exists $T_0 \in \mathbb{N}$ and a sequence of distributions $\mathbb{P}'_T \in \mathcal{P}_T$, $T \in \mathbb{N}$, such that $\mathbb{P}'_T \in \operatorname{int}\mathcal{D}$ for all $T \geq T_0$ and

$$\inf_{\mathbb{P}'\in\operatorname{int}\mathcal{D}} \mathrm{I}(\mathbb{P}', \mathbb{P}) = \liminf_{T\to\infty} \mathrm{I}(\mathbb{P}'_T, \mathbb{P}). \tag{25}$$

Fix any $T \geq T_0$ and let $(i_1, \ldots, i_T)$ be a sequence of observations that generates $\mathbb{P}'_T$. Then, we have

$$\begin{aligned}
\mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{D}) &\geq \mathbb{P}^\infty(\hat{\mathbb{P}}_T = \mathbb{P}'_T) \\
&= |C_T(\mathbb{P}'_T)| \cdot \mathbb{P}^\infty(\xi_1 = i_1, \ldots, \xi_T = i_T) \\
&\geq (T+1)^{-d} \cdot e^{TH(\mathbb{P}'_T)} \cdot e^{T\sum_{i\in\Xi} \mathbb{P}'_T(i)\log\mathbb{P}(i)} \\
&= (T+1)^{-d} e^{-T\mathrm{I}(\mathbb{P}'_T, \mathbb{P})},
\end{aligned}$$

where the first inequality holds because $\mathbb{P}'_T \in \operatorname{int}\mathcal{D} \subseteq \mathcal{D}$, while the second inequality follows from (22) and the lower estimate in (23). This implies that

$$\frac{1}{T}\log\mathbb{P}^\infty(\hat{\mathbb{P}}_T \in \mathcal{D}) \geq -\frac{d\log(T+1)}{T} - \mathrm{I}(\mathbb{P}'_T, \mathbb{P}) \quad \forall T \geq T_0.$$

Taking the limit inferior as T tends to infinity on both sides of the above inequality and using (25) yields the postulated lower bound (7b). This completes the proof. □

Proof of Proposition 4. By construction, the relative entropy $\mathrm{I}(\mathbb{P}', \mathbb{P})$ is twice continuously differentiable in $\mathbb{P}$ at $\mathbb{P} = \mathbb{P}'$. To derive the second-order Taylor expansion around $\mathbb{P}'$, we note that $\mathrm{I}(\mathbb{P}', \mathbb{P}') = 0$ and $\nabla_{\mathbb{P}}\mathrm{I}(\mathbb{P}', \mathbb{P})|_{\mathbb{P}=\mathbb{P}'} = 0$ by Proposition 1(i). Moreover, we have

$$\begin{aligned}
\nabla^2_{\mathbb{P}}\mathrm{I}(\mathbb{P}', \mathbb{P})\big|_{\mathbb{P}=\mathbb{P}'} &= \nabla^2_{\mathbb{P}}\left[H(\mathbb{P}') + \mathrm{I}(\mathbb{P}', \mathbb{P})\right]\big|_{\mathbb{P}=\mathbb{P}'} \\
&= -\frac{1}{T}\,\nabla^2_{\mathbb{P}}\left[\log\mathbb{P}^\infty(\xi_1 = i_1, \ldots, \xi_T = i_T)\right]\big|_{\mathbb{P}=\mathbb{P}'} = F(\mathbb{P}'),
\end{aligned}$$

where the first equality holds because the entropy $H(\mathbb{P}') = -\sum_{i=1}^d \mathbb{P}'(i)\log\mathbb{P}'(i)$ is independent of $\mathbb{P}$, and the second equality follows from a direct calculation; see also (Cover and Thomas 2006, Theorem 11.1.2). Using the definition of $\mathrm{I}(\mathbb{P}', \mathbb{P})$ and the assumption that $\mathbb{P}' > 0$, one further verifies that $\nabla^2_{\mathbb{P}}\mathrm{I}(\mathbb{P}', \mathbb{P})|_{\mathbb{P}=\mathbb{P}'} = \operatorname{diag}(\mathbb{P}')^{-1} \succ 0$. The rest of the proof follows immediately from the definition of the Pearson divergence. □

The following technical lemma is instrumental for the proof of Theorem 9 in the main text.

LEMMA 1 **(A quadratically constrained quadratic program)**. *If $\mathbb{P}' \in \mathcal{P}$, $g \in \Re^d$, $F \in \mathbb{S}^d$ is strictly positive definite, $\mathcal{P}_0 = \{\mathbb{P} \in \Re^d : \sum_{i=1}^d \mathbb{P}(i) = 1\}$ is the affine hull of $\mathcal{P}$, and the columns of $B \in \Re^{d\times(d-1)}$ form a basis for $\{\mathbb{P} - \mathbb{P}' : \mathbb{P} \in \mathcal{P}_0\}$, then*

$$\max_{\mathbb{P}\in\mathcal{P}_0}\left\{g^\top(\mathbb{P}-\mathbb{P}') : \frac{1}{2}(\mathbb{P}-\mathbb{P}')^\top F(\mathbb{P}-\mathbb{P}') \leq r\right\} = \sqrt{2g^\top B(B^\top FB)^{-1}B^\top g}\cdot\sqrt{r}.$$

Proof. Every element of $\{\mathbb{P} - \mathbb{P}' : \mathbb{P} \in \mathcal{P}_0\}$ can be expressed as By for some $y \in \Re^{d-1}$. Thus,

$$\begin{aligned}
\max_{\mathbb{P}\in\mathcal{P}_0}\left\{g^\top(\mathbb{P}-\mathbb{P}') : \frac{1}{2}(\mathbb{P}-\mathbb{P}')^\top F(\mathbb{P}-\mathbb{P}') \leq r\right\} &= \max_{y\in\Re^{d-1}}\left\{g^\top By : \frac{1}{2}y^\top BFBy \leq r\right\} \\
&= \max_{z\in\Re^{d-1}}\left\{g^\top B(B^\top FB)^{-\frac{1}{2}}z : \|z\| \leq \sqrt{2r}\right\} \\
&= \sqrt{2g^\top B(B^\top FB)^{-1}B^\top g}\cdot\sqrt{r},
\end{aligned}$$

where the second equality follows from the substitution $z = (B^\top F B)^{\frac{1}{2}} y$ and the observation that $B^\top F B$ is strictly positive definite, while the third equality holds because the maximization problem over z is solved by

$$z^\star = \sqrt{2r}\, \frac{B(B^\top F B)^{-\frac{1}{2}} B^\top g}{\|B(B^\top F B)^{-\frac{1}{2}} B^\top g\|}.$$

Thus, the claim follows. □

Proof of Theorem 9. Throughout the proof we denote the vector of scenario-wise costs associated with the decision x as $g = (\gamma(x,1), \ldots, \gamma(x,d))^\top$. By Proposition 4, there exists $\lambda > 0$ with $F(\mathbb{P}') \succ \lambda \mathbb{I}$. Moreover, as $\mathbb{P}' > 0$, there exists a sufficiently small $r_0 > 0$ with

$$\mathcal{N} = \left\{ \mathbb{P} \in \Re^d \;:\; \sum_{i=1}^d \mathbb{P}(i) = 1,\; \|\mathbb{P} - \mathbb{P}'\|_2 \leq \sqrt{\frac{2r_0}{\lambda}} \right\} \subseteq \mathcal{P}$$

and $\mathrm{I}(\mathbb{P}', \mathbb{P}) \geq \frac{\lambda}{2} \|\mathbb{P} - \mathbb{P}'\|_2^2$ for any $\mathbb{P} \in \mathcal{N}$. Thus, $\|\mathbb{P} - \mathbb{P}'\|_2 \leq \sqrt{2\mathrm{I}(\mathbb{P}', \mathbb{P})/\lambda}$ for $\mathbb{P} \in \mathcal{N}$.

Proposition 4 further implies that there exist two non-negative functions $\kappa_1, \kappa_2 \in o(1)$ such that

$$\frac{1}{2(1+\kappa_1(r))} \cdot (\mathbb{P} - \mathbb{P}')^\top F(\mathbb{P}')(\mathbb{P} - \mathbb{P}') \;\leq\; \mathrm{I}(\mathbb{P}', \mathbb{P}) \;\leq\; \frac{1+\kappa_1(r)}{2} \cdot (\mathbb{P} - \mathbb{P}')^\top F(\mathbb{P}')(\mathbb{P} - \mathbb{P}') \tag{26a}$$

$$c(x, \mathbb{P}') + g^\top (\mathbb{P} - \mathbb{P}') - \kappa_2(r) \cdot \sqrt{r} \;\leq\; c(x, \mathbb{P}) \;\leq\; c(x, \mathbb{P}') + g^\top (\mathbb{P} - \mathbb{P}') + \kappa_2(r) \cdot \sqrt{r} \tag{26b}$$

for all $\mathbb{P} \in \mathcal{N}$ with $\mathrm{I}(\mathbb{P}', \mathbb{P}) \leq r$ and for all $r \in [0, r_0]$. Indeed, the inequalities (26) hold because $\mathrm{I}(\mathbb{P}', \mathbb{P})$ is locally equal to the quadratic form $\frac{1}{2}(\mathbb{P} - \mathbb{P}')^\top F(\mathbb{P}')(\mathbb{P} - \mathbb{P}')$ and the parametric predictor $c(x, \mathbb{P})$ is locally equal to $c(x, \mathbb{P}') + g^\top (\mathbb{P} - \mathbb{P}')$, while $\mathrm{I}(\mathbb{P}', \mathbb{P}) \leq r$ implies that $\|\mathbb{P} - \mathbb{P}'\|_2 \leq \sqrt{2r/\lambda}$. In the following, let $\mathcal{P}_0 = \{\mathbb{P} \in \Re^d : \sum_{i=1}^d \mathbb{P}(i) = 1\}$ be the affine hull of $\mathcal{P}$, and assume that the columns of $B \in \Re^{d \times (d-1)}$ form a basis for $\{\mathbb{P} - \mathbb{P}' : \mathbb{P} \in \mathcal{P}_0\}$. Then, we have

$$\begin{aligned}
\hat{c}_r(x, \mathbb{P}') \;&\leq \max_{\mathbb{P} \in \mathcal{P}_0} \left\{ c(x, \mathbb{P}') + g^\top (\mathbb{P} - \mathbb{P}') + \kappa_2(r) \cdot \sqrt{r} : \frac{1}{2(1+\kappa_1(r))} \cdot (\mathbb{P} - \mathbb{P}')^\top F(\mathbb{P}')(\mathbb{P} - \mathbb{P}') \leq r \right\} \\
&= c(x, \mathbb{P}') + \frac{1}{\sqrt{1+\kappa_1(r)}} \cdot \sqrt{2 g^\top B (B^\top F(\mathbb{P}') B)^{-1} B^\top g} \cdot \sqrt{r} + \kappa_2(r) \cdot \sqrt{r} \\
&= c(x, \mathbb{P}') + \sqrt{2 g^\top B (B^\top F(\mathbb{P}') B)^{-1} B^\top g} \cdot \sqrt{r} + o(\sqrt{r}),
\end{aligned}$$

where the inequality follows from overestimating the objective function in (10) using (26b) and relaxing the feasible set of (10) using (26a). The first equality relies on Lemma 1, and the second equality exploits the assumption that $\kappa_1, \kappa_2 \in o(1)$. Similarly, we obtain

$$\begin{aligned}
\hat{c}_r(x, \mathbb{P}') \;&\geq \max_{\mathbb{P} \in \mathcal{P}_0} \left\{ c(x, \mathbb{P}') + g^\top (\mathbb{P} - \mathbb{P}') - \kappa_2(r) \cdot \sqrt{r} : \frac{1+\kappa_1(r)}{2} \cdot (\mathbb{P} - \mathbb{P}')^\top F(\mathbb{P}')(\mathbb{P} - \mathbb{P}') \leq r \right\} \\
&= c(x, \mathbb{P}') + \sqrt{1+\kappa_1(r)} \cdot \sqrt{2 g^\top B (B^\top F(\mathbb{P}') B)^{-1} B^\top g} \cdot \sqrt{r} - \kappa_2(r) \cdot \sqrt{r} \\
&= c(x, \mathbb{P}') + \sqrt{2 g^\top B (B^\top F(\mathbb{P}') B)^{-1} B^\top g} \cdot \sqrt{r} + o(\sqrt{r}).
\end{aligned}$$

The above estimates imply that

$$\hat{c}_r(x, \mathbb{P}') = c(x, \mathbb{P}') + \sqrt{2 g^\top B (B^\top F(\mathbb{P}') B)^{-1} B^\top g} \cdot \sqrt{r} + o(\sqrt{r}). \tag{27}$$

In the following we assume without loss of generality that

$$B = \begin{pmatrix} +1 & \cdots & +1 \\ -1 & & \\ & \ddots & \\ & & -1 \end{pmatrix} \in \Re^{d \times (d-1)}.$$

Note that the columns of B form indeed a basis for $\{\mathbb{P}-\mathbb{P}' : \mathbb{P}\in\mathcal{P}_0\}$. Moreover, note that the expansion (27) is unique and thus independent of the particular choice of B. Since $F(\mathbb{P}')=\operatorname{diag}(\mathbb{P}')$ by Proposition 4, we then find via elementary manipulations that

$$B^\top F(\mathbb{P}')B = \operatorname{diag}(\mathbb{P}')^{-1} + \frac{\mathbf{e}\mathbf{e}^\top}{\mathbb{P}'(1)},$$

where $\mathbf{e}$ denotes the vector of ones in $\Re^{d-1}$. The Sherman-Morrison formula further yields

$$(B^\top F(\mathbb{P}')B)^{-1} = \operatorname{diag}(\mathbb{P}'_{-1}) - (\mathbb{P}'_{-1})(\mathbb{P}'_{-1})^\top \quad \text{and} \quad B(B^\top F(\mathbb{P}')B)^{-1}B^\top = \operatorname{diag}(\mathbb{P}') - (\mathbb{P}')(\mathbb{P}')^\top,$$

where $\mathbb{P}_{-1}=(\mathbb{P}'_2,\ldots,\mathbb{P}'_d)^\top$. Recalling the definition of g we finally obtain

$$\begin{aligned} g^\top B(B^\top F(\mathbb{P}')B)^{-1}B^\top g &= \sum_{i=1}^d \mathbb{P}'(i)\gamma(x,i)^2 - \left(\sum_{i=1}^d \mathbb{P}'(i)\gamma(x,i)\right)^2 \\ &= \sum_{i=1}^d \mathbb{P}'(i)\left(\gamma(x,j) - \sum_{j=1}^d \mathbb{P}'(j)\gamma(x,j)\right)^2, \end{aligned}$$

where we use the fact that the variance of a random variable equals the mean of its square minus the square of its mean. The approximation (21b) now follows by substituting the last expression into (27). The equivalence of (21a) and (21b) is a direct consequence of Lemma 1. □

Appendix B: Conic representation of the optimal predictor

We give here an explicit representation of the optimal data-driven predictor $\hat{c}_r(x,\mathbb{P}')$ in terms of a conic optimization problem. For generic $\mathbb{P}'\in\mathcal{P}$ this problem involves the exponential cone. For empirical distributions $\mathbb{P}'\in\mathcal{P}$ based on T samples, however, this problem reduces to a second-order cone program.

Definition 10 (Exponential cone). The exponential cone $K_{\exp}$ is the three dimensional convex cone

$$K_{\exp} = \left\{(x,y,z)\in\Re^3 \;:\; z>0,\;\; \exp(x/z)\le y/z\right\} \cup \left\{(x,y,z)\in\Re^3 \;:\; x\le 0,\;\; y\ge 0,\;\; z=0\right\}.$$

The exponential cone is convex and admits an efficient self-concordant barrier function (Chares 2009). Furthermore, there exists open source optimization routines for solving generic exponential cone programs such as `ECOS-EXP` by Domahidi et al. (2013) and Serrano (2015). The following result is elementary and well-known; see, *e.g.*, Chandrasekaran and Shah (2017) for further references on relative entropy optimization.

Theorem 10 **(Exponential conic representation of $\hat{c}_r$).** *For any distribution $\mathbb{P}'$ in $\mathcal{P}$, the optimal predictor $\hat{c}_r(x,\mathbb{P}')$ can be expressed as the optimal value of the convex optimization problem*

$$\begin{array}{ll} \hat{c}_r(x,\mathbb{P}') = \max & c(x,\mathbb{P}) \\ \text{s.t.} & \mathbb{P}\in\mathcal{P},\;\; \mathbb{Q}\in\mathbb{R}^d_+,\;\; \sum_{i\in\Xi}\mathbb{Q}(i)\le r \\ & (-\mathbb{Q}(i),\mathbb{P}(i),\mathbb{P}'(i))\in K_{\exp} \quad \forall i\in\Xi. \end{array} \tag{28}$$

Proof. By Definition 10 we have

$$\begin{aligned} (-\mathbb{Q}(i),\mathbb{P}(i),\mathbb{P}'(i))\in K_{\exp} &\iff \exp(-\mathbb{Q}(i)/\mathbb{P}'(i))\le \mathbb{P}(i)/\mathbb{P}'(i),\;\; \mathbb{P}'(i)>0 \;\text{ or }\; \mathbb{Q}(i)\ge 0,\;\; \mathbb{P}(i)\ge 0,\;\; \mathbb{P}'(i)=0 \\ &\iff -\mathbb{Q}(i)/\mathbb{P}'(i)\le \log(\mathbb{P}(i)/\mathbb{P}'(i)),\;\; \mathbb{P}'(i)>0 \;\text{ or }\; \mathbb{Q}(i)\ge 0,\;\; \mathbb{P}(i)\ge 0,\;\; \mathbb{P}'(i)=0 \\ &\iff \mathbb{Q}(i)\ge \mathbb{P}'(i)\log(\mathbb{P}'(i)/\mathbb{P}(i)),\;\; \mathbb{P}'(i)>0 \quad\;\; \text{or }\; \mathbb{Q}(i)\ge 0,\;\; \mathbb{P}(i)\ge 0,\;\; \mathbb{P}'(i)=0 \\ &\iff \mathbb{Q}(i)\ge \mathbb{P}'(i)\log(\mathbb{P}'(i)/\mathbb{P}(i)),\;\; \mathbb{P}'(i)\ge 0, \end{aligned}$$

where the last equivalence follows from our convention that $0\log 0=0$. By the definition of the relative entropy, the constraint $\sum_{i\in\Xi}\mathbb{Q}(i)\le r$ thus holds if and only if $\mathrm{I}(\mathbb{P}',\mathbb{P})\le r$. □

Theorem 10 indicates that evaluating the optimal predictor $c_r(x, \mathbb{P}')$ at a generic $\mathbb{P}' \in \mathcal{P}$ reduces to solving a tractable optimization problem involving $\mathcal{O}(d)$ convex exponential cone constraints. Although numerical routines for solving problem (28) exist, most mature or commercial optimization routines require constraints involving self-dual cones. As the exponential cone $K_{\exp}$ is not self-dual, problem (28) is outside of their realm.

In practice the predictor $c_r(x, \mathbb{P}')$ needs to be evaluated only at empirical distributions $\mathbb{P}' \in \mathcal{P}_T = \mathcal{P} \cap [0/T, \ldots, T/T]^d$, where T denotes the number of available samples. Luckily, in this case $\hat{c}_r(x, \mathbb{P}')$ can be represented in terms of a second-order cone program.

THEOREM 11 **(Second-order conic representation of $\hat{c}_r$).** *For any empirical distribution $\mathbb{P} \in \mathcal{P}_T$, the optimal predictor $\hat{c}_r(x, \mathbb{P}')$ coincides with the optimal value of the convex optimization problem*

$$\begin{array}{rll} \hat{c}_r(x, \mathbb{P}') = \max & c(x, \mathbb{P}) & \\ \text{s.t.} & \mathbb{P} \in \mathcal{P} & \\ & \left(\prod_{i \in \Xi} \mathbb{P}(i)^{T \cdot \mathbb{P}'(i)}\right)^{1/T} \geq e^{-(r + H(\mathbb{P}'))}, & \end{array} \tag{29}$$

where $H(\mathbb{P}') = -\sum_{i \in \Xi} \mathbb{P}'(i) \log \mathbb{P}'(i)$ denotes the entropy of the empirical distribution $\mathbb{P}'$.

Proof. The constraint $\mathrm{I}(\mathbb{P}, \mathbb{P}) \leq r$ is equivalent to

$$\begin{aligned} \textstyle\sum_{i \in \Xi} \mathbb{P}'(i) \log(\mathbb{P}'(i)/\mathbb{P}(i)) \leq r &\iff \textstyle\sum_{i \in \Xi} \mathbb{P}'(i)[\log \mathbb{P}'(i) - \log(\mathbb{P}(i))] \leq r \\ &\iff \textstyle\sum_{i \in \Xi} \mathbb{P}'(i) \log(\mathbb{P}(i)) \geq -(r + H(\mathbb{P}')) \\ &\iff \exp(\textstyle\sum_{i \in \Xi} \mathbb{P}'(i) \log(\mathbb{P}(i))) \geq \exp(-(r + H(\mathbb{P}'))) \\ &\iff \textstyle\prod_{i \in \Xi} \exp(\mathbb{P}'(i) \log(\mathbb{P}(i))) \geq \exp(-(r + H(\mathbb{P}'))) \\ &\iff \textstyle\prod_{i \in \Xi} \mathbb{P}(i)^{\mathbb{P}'(i)} \geq \exp(-(r + H(\mathbb{P}'))) \\ &\iff \left(\textstyle\prod_{i \in \Xi} \mathbb{P}(i)^{T \cdot \mathbb{P}'(i)}\right)^{1/T} \geq \exp(-(r + H(\mathbb{P}'))), \end{aligned}$$

and thus the claim follows. □

Note that for empirical distributions $\mathbb{P} \in \mathcal{P}_T$ the quantity $T \cdot \mathbb{P}'(i)$ counts the number of times state i has been observed. The last constraint in (29) thus requires the geometric mean of the vector

$$(\underbrace{\mathbb{P}(1), \ldots, \mathbb{P}(1)}_{T \cdot \mathbb{P}'(1) \text{ times}}, \ldots, \underbrace{\mathbb{P}(d), \ldots, \mathbb{P}(d)}_{T \cdot \mathbb{P}'(d) \text{ times}})$$

of length T to be non-inferior to $e^{-(r + H(\mathbb{P}'))}$. By Nesterov and Nemirovskii (1994, Section 6.2.3.5), this requirement can be re-expressed in terms of $\mathcal{O}(T)$ second-order cone constraints involving $\mathcal{O}(T)$ auxiliary variables. Problem (29) is hence amenable to second-order cone solvers such as `MOSEK`, `GUROBI` or `CPLEX`. Although elementary, to our best knowledge Theorem 11 and the second-order cone representability of $\hat{c}_r(x, \mathbb{P}')$ for empirical distributions $\mathbb{P}' \in \mathcal{P}_T$ are new.